



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification⁶:

G06K 9/00

A1

(11) International Publication Number:

WO 97/21188

(43) International Publication Date:

12 June 1997 (12.06.97)

(21) International Application Number: PCT/US96/19132

(22) International Filing Date: 4 December 1996 (04.12.96)

(30) Priority Data:

60/007,906	4 December 1995 (04.12.95)	US
60/023,037	2 August 1996 (02.08.96)	US
60/027,653	4 October 1996 (04.10.96)	US

(71) Applicant: DAVID SARNOFF RESEARCH CENTER, INC.
[US/US]; 201 Washington Road CN5300, Princeton, NJ
08540-5300 (US).

(72) Inventors: HANNA, Keith, James; 212 John Street, Princeton,
NJ 08540 (US). BURT, Peter, J.; 80 Wheatstheaf Lane,
Princeton, NJ 08540 (US). PELEG, Shmuel; 130 Moore
Street, Princeton, NJ 08540 (US). HERMAN, Joshua, R.;
29 Jared Drive, Robbinsville, NJ 08691 (US). DIXON,
Douglas, F.; 13 Center Street, Hopewell, NJ 08525 (US).
MISHRA, Depan; 102 Fox Run Drive, Plainsboro, NJ 08536
(US). WIXSON, Lambert, E.; 38 Washington Street, Rocky
Hill, NJ 08553 (US). MANDELBAUM, Robert; 1500
Locust Street, Philadelphia, PA 19102 (US). COYLE, Peter;
4 Pansy Court, Newtown, PA 18940 (US).

(74) Agent: BURKE, William, J.; David Saroff Research Center,
Inc., 201 Washington Road CN5300, Princeton, NJ 08543
(US).

(81) Designated States: AM, AT, AU, BB, BG, BR, BY, CA, CH,
CN, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IS, JP, KE,
KG, KP, KR, KZ, LK, LR, LT, LU, LV, MD, MG, MN,
MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI,
SK, TJ, TM, TT, UA, UG, UZ, VN, ARIPO patent (KE,
LS, MW, SD, SZ, UG), European patent (AT, BE, CH, DE,
DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE),
OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR,
NE, SN, TD, TG).


Published

With international search report.

(54) Title: WIDE FIELD OF VIEW/NARROW FIELD OF VIEW RECOGNITION SYSTEM AND METHOD

(57) Abstract

A recognition system which obtains and analyzes images of at least one object in a scene comprising a wide field of view (WFOV) imager (10, 12) which is used to capture an image of a scene and to locate the object and a narrow field of view (NFOV) imager (14) which is responsive to the location information provided by the WFOV imager and which is used to capture an image of the object (56), the image of the object having a higher resolution than the image captured by the WFOV imager is disclosed. In one embodiment, a system that obtains and analyzes images of the irises of eyes of a human or animal in an image with little or no active involvement by the human or animal is disclosed. A method of obtaining and analyzing the images of at least one object in a scene comprising capturing a wide field of view image of the object to locate the object in the scene; and then using a narrow field of view imager responsive to the location information (54) provided in the capturing step to obtain a higher resolution image of the object is also disclosed.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic			SE	Sweden
CG	Congo	KR	Republic of Korea	SG	Singapore
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LR	Liberia	SZ	Swaziland
CS	Czechoslovakia	LT	Lithuania	TD	Chad
CZ	Czech Republic	LU	Luxembourg	TG	Togo
DE	Germany	LV	Latvia	TJ	Tajikistan
DK	Denmark	MC	Monaco	TT	Trinidad and Tobago
EE	Estonia	MD	Republic of Moldova	UA	Ukraine
ES	Spain	MG	Madagascar	UG	Uganda
FI	Finland	ML	Mali	US	United States of America
FR	France	MN	Mongolia	UZ	Uzbekistan
GA	Gabon	MR	Mauritania	VN	Viet Nam

WIDE FIELD OF VIEW/ NARROW FIELD OF VIEW RECOGNITION SYSTEM AND METHOD

The invention is directed to video image capture and processing systems and
5 methods therefor.

There are many methods for recognizing or identifying an individual on a
transactional basis. These include analyzing a signature, obtaining and analyzing an
image of a fingerprint, and imaging and analyzing the retinal vascular patterns of a
human eye. All of these recognition techniques have a common drawback, they require
10 the individuals being recognized to perform some positive act, either signing their
names, placing one of their fingers on an imaging plate or positioning themselves using
a head-rest or bite bar or position their eye relative to an eyepiece so that an image of
their vascular patterns may be captured.

Recently iris capture and analysis has been gaining favor as a method for
15 identifying individuals. This technique is described in U.S. Patent Nos. 4,641,349.
and 5,291,560 and in an article by J.G. Daugman in the *IEEE Transactions on Pattern
Analysis and Machine Intelligence*, vol. 15, no. 11, pp. 1148-61, November 1993.
The systems described in these references require the person being identified to hold at
least one of their eyes in a fixed position with respect to an imaging imager while their
20 iris is being imaged. While this procedure may be satisfactory for some applications, it
is not satisfactory for quick transactional activities such as using an automatic teller
machine (ATM). A person using an ATM typically does not want to spend any more
time than is necessary to complete his or her transactions. Consequently, it would be
considered an undue burden to ask ATM users to perform any positive act other than
25 inserting their ATM card and keying in their desired transactions on a keypad.

This need makes clear that there exists a more general problem of identifying
objects or individuals in a passive way that is both fast and accurate.

SUMMARY OF THE INVENTION

The invention is embodied in a system which obtains and analyzes images of at
30 least one object in a scene comprising a wide field of view (WFOV) imager which is
used to capture an image of the scene and to locate the object and a narrow field of view
(NFOV) imager which is responsive to the location information provided by the WFOV
imager and which is used to capture an image of the object, the image of the object
having a higher resolution than the image captured by the WFOV imager.

35 The invention is embodied in an automatic system that obtains and analyzes
images of the irises of eyes of a human or animal in an image with little or no active
involvement by the human or animal. According to one aspect of the invention, the
system includes both WFOV and NFOV imagers. The system includes control

circuitry which obtains an image from the WFOV imager to determine the location of the eyes and then uses the NFOV imager to obtain a high-quality image of one or both of the irises.

The invention is also a method for obtaining and analyzing images of at least one object in a scene comprising capturing a wide field of view image of the object to locate the object in the scene; and then using a narrow field of view imager responsive to the location information provided in the capturing step to obtain higher resolution image of the object.

BRIEF DESCRIPTION OF THE DRAWING

Figure 1a is a block diagram of an acquisition system according to one embodiment of the invention.

Figure 1b is a flow chart useful for describing the operation of the acquisition system of Figure 1a.

Figure 1c is a front plan view of an ATM which includes an embodiment of the invention.

Figure 2 is a front plan view of a physically smaller iris recognition system.

Figure 3 is a functional block diagram of apparatus suitable for use in the ATM of Figure 1c.

Figures 4a, 4b, 5a, 5b, 5c, 5d and 5e illustrate alternative configurations of the light source, NFOV imager and pan and tilt mirror for the apparatus shown in Figures 1 and 3.

Figures 6, 7, 8 and 9 are drawings of a person that are useful for describing the operation of the apparatus shown in Figures 1c and 3.

Figures 10, 11, 12 and 13 are drawings representing a human eye that are useful for describing the operation of the apparatus shown in Figures 1c and 3.

Figure 14 is a flow-chart illustrating the high-level control flow for the control processor shown in Figure 3.

Figure 15 is a flow-chart illustrating a process that details the process step in Figure 14 which locates the head and eyes of the individual.

Figure 16 is a flow-chart illustrating details of the process step in Figure 15 which locates the head in the image.

Figure 17 is a flow-chart illustrating an implementation of the process step in Figure 15 which identifies possible facial features.

Figure 18a is a flow-chart illustrating an implementation of the symmetry analysis block shown in Figure 15.

Figures 18b, 18c, 18e through 18k and 18m, 18n, 18r, and 18v are flow-chart diagrams which illustrate an alternative method for locating the user's eyes in the WFOV images.

Figure 18d is a drawing representing a human head that is useful for describing the exemplary method for locating the user's eyes.

Figure 18l is a drawing of a cone shaped search region for the specularly process.

5 Figures 18o -18q and 18s - 18u are diagrams useful for explaining the specularly detection process.

Figure 19 is a flow-chart illustrating a method of implementing the find range block shown in Figure 14.

10 Figure 20 is a flow-chart illustrating a method of implementing the locate iris block of the flow-chart shown in Figure 14.

Figure 21 is a flow-chart illustrating a method of implementing the obtain high quality image block of the flow-chart shown in Figure 14.

Figure 21a is a flow chart for describing the detection of specularities in the NFOV imagery.

15 Figure 21b is a diagram illustrating the process of Figure 21.

Figure 22a is a flow-chart of a method of implementing the circle finder step of the flow-chart shown in Figures 20 and 21.

Figure 22b is a drawing of a human eye which is useful for describing the operation of the method of Figure 22a.

20 Figure 23 is a flow-chart of a method for implementing the extract iris step of the flow-chart shown in Figure 21.

Figure 24 is a flow-chart of a method for locating the person who is to be identified using iris recognition and determining the distance that person is from the acquisition system.

25 Figure 24a is a flow-chart of a method for producing a region of interest (ROI containing the user's head) for the step 1420 shown in Figure 24.

Figure 25 is a flow-chart of a method for using the depth acquired using the method shown in Figure 24 to adjust the NFOV imager on the user.

30 Figure 26 is a flow-chart of a method for calibrating the system and generating the values stored in the LUT described in Figure 25.

Figure 27 is a flow-chart illustrating a method of autofocus for the NFOV imager.

35 Figure 28 is a flow-chart of a method for detecting the user's eyes using reflection off the back of the eye and the occluding properties of the iris/pupil boundary.

Figure 29 is a flow-chart of a method of removing ambient specular reflections from an image.

Figure 30 is a diagram of the mounting apparatus for the WFOV imagers.

Figure 31 is a flow-chart of a method for adjusting the mounting apparatus shown in Figure 30.

Figure 32 is a block diagram of a test setup used in the process shown in Figure 31 for adjusting the mounting apparatus.

Figures 33 and 34 are perspective views of an another embodiment for detecting barcodes using WFOV imagery and NFOV imagery.

Figure 35a is a perspective view of the barcodes on a container.

Figure 35b and 35c are exemplary barcodes for use with the system shown in Figures 33 and 34.

DETAILED DESCRIPTION

While the invention is described in terms of an iris recognition system which verifies the identify of an person, it is contemplated that it may be practiced more generally as a system and method which uses successive WFOV processing and NFOV processing to locate and identify objects in a scene which contains the objects.

The exemplary embodiment of the invention is directed to an automated acquisition system for the non-intrusive acquisition of images of human irises for the purpose of identity verification. This embodiment uses active machine vision techniques that do not require the user to make physical contact with the system, or to assume a particular pose except that the user preferably stands with his head within a designated calibrated volume.

In Figure 1a, the system 5 consists of a stereo pair of wide field-of-view (WFOV) imagers 10 and 12, such as video imagers, a narrow field-of-view (NFOV) imager 14, such as a video imager, a pan-tilt mirror 16 allowing the image area of the NFOV imager to be moved relative to the WFOV imagers 10 and 12, an image processor 18 which may be a PV-1™ real-time vision computer available from Sensar Inc. or David Sarnoff Research Center, and a processor 20 which may be any of a number of computers which use a PENTIUM™ microprocessor manufactured from Intel, Corp., or other microprocessors having similar capabilities. The system 5 actively finds the position of a user's eye 30 and acquires a high-resolution image to be processed by the image processor 18 which performs iris recognition.

The operation of the system 5 is described below with reference to Figure 1b. The head and depth finding process 50 uses a pair of stereo WFOV images from WFOV imagers 10 and 12. Using the stereo images, process 50 selects the nearest user to the system, finds the position of the user's head in the image, and estimates the depth of the user's eye 30 from the system 5. Process 50 implements a cross-correlation-based stereo algorithm to build a disparity map of the WFOV scene, the scene acquired by the WFOV imagers 10 and 12. The disparity map is then analyzed and the closest region, the region of interest (ROI), of approximately the size and shape

corresponding to that of a human head is extracted. The disparity corresponding to the user's face is then taken to be the mean disparity of this region. The three dimensional depth of the user's face is proportional to the inverse of the disparity.

An image of the ROI containing the user's head is provided to WFOV eye finding process 52 which locates the user's eye is located in the ROI. It is also contemplated that the estimated depth of the user's head may also be provided to the WFOV eye finding process 52. Process 52, may, for example, use a template to locate the user's right eye in the ROI. Alternatively, the user's right or left eye could be located using one of the other processes described below. Process 52 may also analyze and combine the results of three eye finding processes to verify and determines the precise location of the user's eye.

The first process is a template based process which locates the user's face in the ROI by searching for characteristic arrangements of features in the face. A band-pass filtered version of the ROI and the orientation of particular features of the user's face, for example, the mouth, at a coarse resolution of the ROI are compared to the template. The face template comprises a priori estimate of the expected spatial arrangement of the facial features. A face is detected when a set of tests using these features is successfully passed.

The second process is a template-based method which uses similar features to those used in the first process but locates the eyes by identifying the specularities from the surface of spectacles worn by the user, if present. The third process is a specularities-based process that locates reflections of the illuminators that are visible on the user's cornea. The information from the first, second, and third process are combined to determine whether an eye has been detected, and, if so, its location in the image.

Next, process 54 maps the depth and WFOV image coordinates of the user's eye to estimate the pan, tilt, and focus parameters of the pan-tilt mirror 16 and the NFOV imager 14 which are used to capture an image of the user's eye with the NFOV imager 14. A calibration look-up table (LUT) is used to map the information recovered from the WFOV processing of processes 50 and 52 onto the NFOV imaging parameters which are used to align the WFOV and NFOV images. The input values to the LUT are the x, y image coordinates of the eye and the depth z of the head. The LUT provides as output values the pan and tilt angles of the pan/tilt mirror 16, the focus of the NFOV imager, and the expected diameter of the iris. The values stored in the LUT account for the baseline separations between the WFOV imagers and the NFOV imager 14, lens distortion in the imagers 10 and 12, and vergence of the imagers 10 and 12.

The contents of the LUT are obtained using a calibration process. Calibration is performed when the acquisition system is built. An object of known size is placed at an

x, y location (relative to the imager) in the image. The depth z in that region is computed by process 50. The pan/tilt mirror is manually slewed so that the object is centered in the NFOV image. The image is then focused, and the diameter in pixels of the object is measured. Thus, for each point, the set of corresponding values {x, z, pan, tilt, focus, iris diameter} is recorded. The x, y, and z values and the three dimensional coordinates of the user's head with respect to the acquisition system 5. Pan and tilt values are the adjustments for the pan-tilt mirror 16. The focus value is the focus of NFOV imager. Finally, iris diameter value is the expected size of the user's iris.

This process is repeated for up to twenty points per depth plane, and at up to four depths inside the working volume of the acquisition system. Next, for each set of neighboring points within the acquired points, a vector of linear functions is fit to the data as shown in relation (1) below.

$$(f_{\text{pan}}, f_{\text{tilt}}, f_{\text{focus}}, f_{\text{diam}}) : X \times Y \times Z \nrightarrow P \times T \times F \times D \quad (1)$$

X is x, Y is y, Z is z, P is pan, T is tilt, F is focus, and D is iris diameter. The result is a set of vectors of functions that define the mapping from X, Y, and Z to each of the described parameters throughout the volume. In operation, the LUT maps the user's eye location (x, y, z) to the linearly interpolated NFOV imager parameters $(f_{\text{pan}}, f_{\text{tilt}}, f_{\text{focus}}, f_{\text{diam}})(x, y, z)$ derived from the values stored in the LUT. Other values may also be stored in the LUT including aperture size, expected specular size, and expected distance between specularities. The aperture size may be determined by placing a white object at each position during calibration and adjusts the operation of the imager so that a uniform brightness level may be established. The expected specular size and the expanded distance between specularities may be used to identify false detection of specularities, that is, specularities that are not from the user's cornea or glasses.

Next, the user's eye 30 is found in the NFOV image using the eye finding process 56. Process 56 detects a set of features visible in the NFOV image if the eye is in the field of view. Process 56 uses two incandescent light sources, one on either side of the NFOV imager. Light from these sources is reflected from the cornea of the eye 30, and appears as two bright spots. The specularities are used both to confirm the presence of the eye in the NFOV image and subsequently to determine the location of the eye 30. The separation of the detected specularities is estimated from the depth information obtained by process 50. Because the specularities are approximately symmetrically located on either side of the eye center, their positions are used to estimate the coordinates of the center of the iris. Once the presence of the eye has been reliably determined, closed-loop NFOV tracking is used without information from the

WFOV image. In the event of large motion by the user, the NFOV imager 14 may lose track of the eye 30. In this instance, process 50 or processes 50 and 52 may be initiated to reacquire the eye 30.

Process 58 checks and adjusts, if necessary, the image quality of the image of the iris which is about 200 by 200 pixels. Image quality is adjusted by electronically centering the eye 30 in the image center from the imager 14 using the mean position of the detected specularities after mechanical centering of the eye in the imager 14. The image of the user's eye is then processed to identify the user.

In addition, the system may store other attributes of the user such a height, facial features, hair color, or face color for recognizing the user. This information may be acquired and stored in the system during enrollment when information relating to the user's iris is acquired and stored. Further, the system may also include security features to ensure that the acquired image of the iris is from a real person and not a imitation. For this purpose, the system may include blink detection processes to detect blinking of the user's eye lid, a pupil size process to detect changes in the user's pupil in response to changes in illumination, or a tremor detection process to detect the natural tremor of a person's eye.

The same components and processes may be used in an enrollment process to store iris information. During the enrollment process, the system would perform the same operations as described above and below except that the acquired image of the user's eye would be stored in a database.

Figure 1c is a front plan view of an ATM which includes an iris recognition system of the invention. In addition to the basic iris recognition system, the ATM shown in Figure 1c also illustrates several alternative illumination schemes.

The ATM includes several features common to all ATM's, a display 110, a keypad 112 and a card reader 134. Although used for a different purpose, most ATM's also include a WFOV imager, such as the imager 10 shown in Figure 1c. In a conventional system, the imager 10 is used to obtain images of persons using the ATM for security purposes. In an ATM of the invention, the WFOV imager is also used to capture an image of a person who is using the ATM in order locate the person's eyes for a subsequent iris imaging operation.

In addition to the WFOV imager 10, the ATM includes a second WFOV imager 12 for stereoscopic imaging, a NFOV imager 14 and a pan and tilt mirror 16. One or both WFOV imagers 10 and 12 may include on-axis illumination 12s and off-axis illumination 12r that may be used in conjunction with the imagers 10 and 12 to detect and remove unwanted specularities from the imagery. The mirror 16 is used to direct light reflected from the person using the ATM into the lens system of the imager 14. The ATM shown in Figure 1c also includes a sonic rangefinder transducer 114 which

may be used, as described below, to determine an approximate distance between the imager 14 and the person using the ATM.

While, in this embodiment, the mirror is on a pan and tilt mounting, it is contemplated that a similar effect could be generated by using a fixed mirror and having the imager 14 on a pan and tilt mounting.

The ATM shown in Figure 1c includes several alternative light sources that are used to illuminate the person for the imagers 10, 12 and 14. The light source 124 is positioned close to the optical axis of the WFOV imager 10. This light source may be used to locate the eyes of the person quickly using the "red eye" effect, as described below with reference to Figures 12, 13 and 15. One alternative lighting scheme includes the light sources 126 and 128. These light sources are positioned close to the imagers 10, 12 and 14 but far enough from the optical axes of the imagers such that a "red-eye" effect is not produced. These light sources may also have different shapes so that their specular reflections in the person's eyes can be differentiated from other reflections used to locate the eyes and to determine their gaze direction, as described below with reference to Figures 11 and 21. The third alternative light sources are the lights 130 and 132 which are located distant from the imagers 10, 12 and 14 and are relatively large in size so as to provide a diffuse illumination.

The ATM system in Figure 1c includes another alternative light source (not shown) which is directed through the mirror 16. This light source is primarily used with the NFOV imager but may be used by the imagers 10 and 12 in much the same way as the light sources 126 and 128, described below with reference to Figure 21.

The iris recognition system may occupy a relatively large area in the ATM depending upon the illumination method that is used. Figure 2 shows a minimal configuration which includes WFOV imagers 10 and 12, a NFOV imager 14 and a pan and tilt mirror 16. In this embodiment, the WFOV imagers use a combination of ambient illumination and a light source (not shown) internal to the ATM which provides light via the mirror 16. In addition, by placing the WFOV imager 10 on axis with the other WFOV imager 12 in the vertical direction errors in determining the position of the user's eyes are minimized. Further, one of the WFOV imagers 10 and 12 may also be arranged to have the same optical axis as the NFOV imager 14. Thus, errors in pointing and focusing the NFOV imager 14 based on the WFOV images may be minimized because the imager 14 is already aligned with one WFOV imager.

In the exemplary iris recognition system in Figure 3, the WFOV driver 312 and the NFOV driver 320 are implemented using the Smart Video Recorder Pro, manufactured by Intel. In addition, the sonic rangefinder 332 is implemented using the SonarRanger Proximity Subsystem Manufactured by Transition Research Corporation. It is contemplated, however that all the drivers and image processors may be

implemented in software on a workstation computer. The host computer on which the control process 310 is implemented, is a Dimension XPS P120c PC computer manufactured by Dell.

In Figure 3, the imagers 10 and 12 are coupled to the WFOV driver 312. A imager suitable for use as either of the imagers 10 or 12 is the IK-M27A Compact Desktop Color Video Camera manufactured by Toshiba. While the imagers are described in terms of video cameras it is understood that the cameras may operate at other than video rates and that any transducer for converting an image of an object or scene into an electrical signal may be used.

The imagers 10 and 12 are both mounted, as shown in Figure 30, on a imager mounting bracket 5000 using their respective tripod mounts 10a and 12a. By adjusting each of the set screws 5010, the imagers 10 and 12 may be moved through several degrees of freedom to align the WFOV imagers 10 and 12. The alignment of the imagers is explained below with reference to Figures 31 and 32.

As shown in Figure 32, the WFOV imagers 10 and 12 are aligned using a target 5080 and display device 5085. The imagers 10 and 12 are aligned so that the images acquired by each WFOV imager are aligned on the display device 5085. In Figure 31, at step 5062 one of the WFOV imagers 10 and 12 is adjusted by adjusting the set screws 5010 to acquire an image of the target 5080 to be displayed, for example, in approximately the center of the display device 5085. At step 5064, the other one of the WFOV imagers 10 and 12 is adjusted to acquire an image of the target 5080 which is also displayed on the display device 5085 and which is aligned with the displayed target of the other WFOV imager.

Then, at step 5068, the alignments of the imagers 10 and 12 are checked using the target 5080 and the display device 5085 and, if necessary, adjusted again using the set screws 5010. At step 5070, holes are drilled through the imager mounting bracket 5000 and into the tripod mounts 10a and 12a and holding pins (not shown), such as split pins, are inserted into the holes. This freezes the WFOV imagers 10 and 12 with respect to the camera mounting bracket 5000 and prevents movement if screws 5040 loosen and loss of alignment. A split pin is a pin that has a tubular shape folded around itself and made round. The holding pins may also be solid pins. Finally, at step 5072, an epoxy is injected between the camera mounting bracket 5000 and the tripod mounts 10a and 12a to further prevent movement. The alignment process described with reference to Figures 31 and 32 is a cost efficient method of aligning the WFOV imagers 10 and 12 which allows rapid alignment of the imagers 10 and 12. The alignment process also enables the mounting hardware to be compact.

The camera mounting bracket 5000 is mounted to a mounting bracket 5030 through orthogonal slots 5031 and 5032. The mounting bracket 5030 is coupled to the

system through slots 5060. Slots 5002, 5032, and 5060 provide movement of the WFOV imagers 10 and 12 relative to the acquisition system so that the pan and tilt of the WFOV imagers 10 and 12 may be adjusted during the manufacture of the ATM.

In Figure 3, the driver 312 obtains images from one or both of the imagers 10 and 12 and provides these images to the host processor at a rate of three to five images per second. The inventors recognize, however, that it would be desirable to have a driver 312 and host interface which can provide image data at a greater rate. The WFOV images are passed by the driver 312 to the stereo face detection and tracking and eye localization module (the stereo module) 316. The stereo module 316 locates portions of the image which include features, such as skin tones or inter-image motion, that may be used by the stereo module 316 to find the person's head and eyes. The locations of the head and eyes determined by the stereo module 316 for each frame of the WFOV image are stored in an internal database 318 along with other, collateral information found by the stereo module 316, such as approximate height, hair color and facial shape. The process 312 and the stereo module 316 are controlled by a control process 310. The stereo module 316 provides a signal to the process 310 when it has located the person's eyes in the image. Further, control information may be provided to the WFOV imagers 10 and 12 via the WFOV driver 312 to control the aperture, focus, and zoom features of the WFOV imagers 10 and 12 that may be stored in a look-up table as described below.

As shown in Figure 3, the WFOV driver 312 is also coupled to receive images from a second WFOV imager 12. Together, the imagers 10 and 12 provide a stereoscopic view of the person using the ATM. Using this stereoscopic image, the stereo module 316 can determine the position of the person's eyes in space, that is to say, it can determine the coordinates of the eyes in an (X, Y, Z) coordinate system. Knowing Z coordinate information about the eyes is useful for focusing the NFOV imager in order to quickly capture a high-quality image of at least one of the person's eyes.

A similar result may be obtained using a single imager and two light sources. Using this technique, the imager 10 obtains two successive images of the person: a first image illuminated only from the left by light source 130 shown in Figure 1c and a second image illuminated only from the right by light source 132. Together, these images provide photometric stereoscopic information about the person. These photometric stereo images may be analyzed in much the same way as the true stereo images in order to determine the distance of the person's eyes from the NFOV imager (i.e. Z coordinate information) as well as the location of the eyes in the WFOV image (i.e. X, Y coordinate information).

Where only a single imager 10 and a single light source are used by the iris recognition system, information about the distance of the person from the NFOV imager can be obtained from the sonic rangefinder 332 and sonic transducer 114. The rangefinder 332 is controlled by the control process 310 to determine the distance
5 between the ATM and the person using conventional ultrasonic ranging techniques. The distance value returned by the rangefinder 332 is passed through the control process 310 to the internal database 318.

Another method of determining Z coordinate distance when only a single imager and a single light source are used is to scan the NFOV imager along a line in the X, Y
10 coordinate plane that is determined by the X, Y coordinate position of the eyes determined from processing the WFOV image. This line corresponds to all possible depths that an image having the determined X, Y coordinates may have. As the NFOV imager is scanned along this line, the images it returns are processed to recognize eye-like features. When an eye is located, the position on the line determines the distance
15 between the near field of view imager and the customer's eyes.

The NFOV imager 14120 is coupled to a NFOV / medium field of view driver 320. The driver 320 controls the focus and zoom of the imager 14 via a control signal F/Z. In addition, the driver 320 controls the mirror 16. A imager suitable for use as the NFOV imager 14 is the EVI-320 Color Camera 2X Telephoto manufactured by
20 Sony and a pan and tilt mirror is the PTU-45 Computer Controlled Pan-Tilt Mirror Adapter manufactured by Directed Perception. In addition, the exemplary NFOV imager 14 uses a 46 mm FA 2X Telephoto Video Converter, manufactured by Phoenix as its zoom lens. In Figures 1, 2 and 3, the imager 14 and its zoom lens (not shown) are mounted in a fixed position in the ATM and the mirror 16 is used to scan image
25 captured by the NFOV imager in the X and Y directions. The focus control on the imager lens is activated by the signal F/Z to scan the imager in the Z coordinate direction. In this embodiment, the zoom control is not used. It is contemplated, however, that the zoom control may be used 1) to magnify or reduce the size of the eye imaged by the near field of view imager in order to normalize the image of the iris or 2)
30 to capture images at a medium field of view (small zoom ratio) prior to capturing images at a NFOV (large zoom ratio).

In Figure 3, the image captured by the NFOV imager 14 is one in which the person's iris has a width of approximately 200 pixels in the high resolution image. In addition, the NFOV driver 320 may capture several images of the eye in close time
35 sequence and average these images to produce the image which is passed to the iris preprocessor 324. This averaged image has reduced noise compared to a single image and provides better feature definition for darkly pigmented irises. Alternatively, other techniques may be used for combining images such as taking the median or a mosaic

process in which several images of marginal quality are combined to form an image of sufficient quality of the iris. This image is also passed to a NFOV/ WFOV image tracking process 322. The driver 320 controls the imager 14 to provide NFOV images to the host processor at a rate of three to five per second although a higher frame rate would be desirable.

The preprocessor 324 locates the boundaries of the iris and separates the portion of the image which corresponds to the iris from the rest of the image returned by the driver 120. This image is normalized to compensate for tilt introduced by the mirror 16, and to compensate for the person having a gaze direction which is not directly into the lens of the NFOV imager. This process is described below with reference to Figures 21 through 23.

When an image of a user wearing glasses is being acquired, it may be desirable for the user to look to one side. The desired angle of the user's head depends on the glass geometry of the glasses. There is a variability of the tilt of the glass surface more about a horizontal axis compared to any other. Therefore, viewing direction with regard to this factor is to the left or the right of the mirror/light source. A second factor is the nose which can occlude illumination. Therefore, the user should be guided to look to the left if the right eye is being imaged, and vice versa.

The separated iris image produced by the preprocessor 324 is passed to the internal database 318 and to an iris classification and comparison process 326. The process 326 receives image information on the person who is using the ATM from a customer database 328. The record in the database corresponding to the customer is identified from data on the ATM card that the person inserted into the card reader 134. Alternatively, it is contemplated that the card itself may be programmed with the customer's iris data. This data may be held in a conventional magnetic stripe on the back of the card or in read-only memory internal to the card if the ATM card is a conventional memory card or "smart" card. In this alternative embodiment, the customer database 328 may hold only the iris data retrieved from the ATM card. This implementation may need an additional data path (not shown) from the card reader 134 to the customer database 328. This path may be implemented via a direct connection or through the user interface process 334 and control process 310.

The image tracking process 322 receives successive NFOV images from the driver 320. Using these images, it correlates facial features from one image to the next and controls the mirror 16, through the driver 320, to keep the iris approximately in the center of the NFOV image. In addition, in this embodiment, the image provided by the driver 320 is 640 pixels by 480 pixels which is less than the 768 pixel by 494 pixel image provided by the imager 14. The driver 320 selectively crops the image returned by the imager 14 to center the iris in the image. Thus, the tracking circuit 322 controls

the mirror 16 and indicates to the driver 320 which portion of the NFOV image is to be cropped in order to keep the user's iris centered in the images returned by the driver 320.

Image tracking based only on the NFOV image is necessarily limited. It can only track relatively small motions or larger motions only if they occur at relatively slow speeds. To augment the image tracking capability of the system, the tracking circuit 322 also receives feature location information from the WFOV image, as provided by the stereo module 316 to the database 318.

Image tracking using WFOV images may be accomplished using a cross correlation technique. Briefly, after image of the head has been located in the WFOV image, it is copied and that copy is correlated to each successive WFOV image that is obtained. As the customer moves, the image of the head moves and the correlation tracks that motion. Further details of this and other image tracking methods are shown in Figure 37 and disclosed in U.S. patent no. 5,063,603, which is hereby incorporated by reference for its teachings on image tracking.

In the exemplary embodiment, the WFOV imager 10, the NFOV imager 14 and the 16 are calibrated by the control process 310 such that features in the WFOV image can be captured in a NFOV image without excessive scanning of the NFOV image in any of the three coordinate directions. This calibration is performed to program the look-up table as described below.

The iris classification and comparison process 326 compares the image to an iris image of the person obtained from the customer database. In the embodiment, two iris images, one for each eye, are held in the database for each customer. The process 326 compares the image returned by the preprocessor 324 to each of these stored images and notifies the control process 310 if a match has been found. In other applications of the iris recognition system, it may be desirable to match the obtained iris image to one of several images held in the customer database or in a similar database. For these uses, the process 326 may classify the iris image using a hash function and then compare the image to only those images which are in the same hash class.

Illumination of the scene being imaged is achieved by the light sources 331 and 321 responsive to the lighting control process 330. The control process 330 may, for example, switch a specified one of the light sources 124, 126, 128, 130 and 132 (collectively light source 331) and the light source 321 on or off and may also control the brightness of any of these light sources. The process 330 is coupled to the control process 310. The process 310 provides the process 330 with specific commands to control the various light sources. It is contemplated, however, that the process 330 may be programmed with sequences of commands such that a single command from process 310 may cause the process 330 to execute a sequence of illumination

operations. In addition, any of the light sources 124, 126, 128, 130, 132 and 321 may be augmented by an infrared light source (not separately shown) such as the TC8245IR Indoor IR Light manufactured by Burle Industries.

In addition, it is contemplated that any or all of the light sources may include a polarization filter and that the imagers 10, 12 and 14 may have a opposite-phase polarization filter. The inventors have determined that circular polarization is more desirable than linear polarization because linearly polarized light may create image artifacts in the iris. This type of polarization greatly reduces specular reflections in the image. These reflections may be, for example, from the customer's corrective lenses. This polarization may be controlled so that the imagers have the same polarization as the light sources when it is desirable to capture images having specular reflections.

Furthermore, it is contemplated that any of the light sources may be a colored light source. This is especially appropriate for the light sources 126 and 128 which are used to produce specular reflections on the customer's eyes.

The control process 310 is also coupled to the main control and user interface process 334, which controls customer interaction with the ATM via the card reader 134, keypad 112 and display 110. In addition, the ATM may include a touch screen 336 (not shown in Figure 1c) through which the customer may indicate selections. The selections made using the touch screen 336, or keypad 112 may be routed to the control process 310 via the main control and user interface process 334. The control process may also communicate with the user via the display 110 through the process 334. In some of the embodiments, for example, it may be desirable to ask the user to stand at a certain minimum distance from the ATM or look at a particular location on the ATM in order to properly capture an image of his or her iris. Implementation of the communication functions between the control process 310 and the display 110 depend on the ATM. These functions could be readily implemented by a person skilled in the art of designing or programming ATMs.

Figures 4a, 4b, 5a and 5b illustrate two alternative physical configurations for the light source 321, mirror 16 and NFOV imager 14. In Figure 4a, the light source 321 is located below the imager 14 to produce a light beam which is at an angle q from the optical axis of the imager. All of the elements are mounted on a platform 410, internal to the ATM. The angle q is selected to be as small as possible and yet produce minimal "red eye" effect. The angle q is approximately 10 degrees. Figures 5a and 5b show an alternative configuration in which two sources 321 are mounted adjacent to the imager 14 on the platform 410. In this implementation, the light sources are also mounted to produce light beams at an angle q from the optical axis of the imager.

Figures 5d and 5e illustrate alternative arrangements of the light source and imager. As is shown in Figure 5d, light 321 provides on-axis illumination with imager

14 using a reflective mirror as is known in the art. Figure 5e shows an embodiment where the light from the light source 321 does not pass through the pan-tilt mirror 122.

Another possible configuration of the light source 321 and the imager 14 is to direct the light from the source 321 to the mirror 16 through a half-silvered mirror (not shown) which is located along the optical axis of the imager 14. This configuration would direct the light along the optical axis of the imager 14. This method of illuminating the customer is used either with a relatively dim light source 321 so as to not produce significant "red eye" effect or in imaging operations in which the "red eye" effect can be tolerated. In this embodiment, if a light source generates illumination which is coaxial with NFOV imager light path, then the light generated by the light source can be steered using the same process as steering the light path for the NFOV imager. More efficient illumination such as reliable LED based illuminators can be used rather than powerful but unreliable incandescent lights which may be used if the whole scene is to be irradiated with, for example, IR light.

When IR is used for illumination, the IR cut off filter of imager 14 is removed. Imager 14 includes an IR cut off filter when IR light is not used by the system. When light sources 321 only generate IR light and visible light is not used, an IR pass filter 123 is positioned in front of the imager 14 as shown in Figure 5c. If both visible light and IR light illumination are used, IR pass filter 123 is not placed in front of imager 14. Further, mirror 16 reflects IR light when IR light is used.

Ambient IR light can cause specular reflections and specular images reflected off the cornea which occlude or corrupt the image of the iris. Thus, it is desirable to remove the specular reflections and specular images caused by the IR light. The light source 321 will not change these specular images, but will increase the amount of light reflected off the iris. The ambient IR light is removed using the process shown in Figure 29. At least one image is acquired using the NFOV imager 14 (shown in Figure 1c) with the light source 321 turned off (shown in Figure 1c) and one image is acquired using the NFOV imager 14 with the light source 321 turned on. The two acquired images are spatially aligned and compared using, for example, image subtraction. The resulting image is the result of light source 321. This assumes that light source 321 and ambient illumination together are within the dynamic range of the NFOV imager 14 so that the resulting image after the image comparison is within the range of the NFOV imager 14.

At step 3110, shown in Figure 29, a first image is acquired using ambient illumination. At step 3120, a second image is acquired using NFOV illumination 1/30 second later. At step 3130, the position of the eye is acquired from the first image using a spoke filter, described below with reference to Figures 21, 22a, and 22b, to identify the iris, such as that described with reference to Figure 22a. At step 3140, the

position of the eye in the second image is acquired using the same process as in step 3130. The two images are coarsely aligned using the identified positions of the eyes in each image. The first image and the second image are precisely aligned using a gradient based motion algorithm. An exemplary gradient based motion algorithm is described in J. R. Bergen, "Hierarchical Model-based Motion Estimation", Proceedings of European Conference on Computer Vision-92, pp. 1-21 (March 23, 1992), which is herein incorporated by reference. At step 3170, the aligned first image and the aligned second image are subtracted to remove specular reflections from the image caused by the IR light.

It may be desirable to illuminate with a combination of light sources including IR light. Existing ambient light can be used as well as IR light. Two images could be encoded on enrollment - one in visible light and one in IR. For iris recognition, an iterative algorithm searches the first image for a poor iris match, finally iterating onto the iris code that corresponds to the best match for a ratio of IR to visible light. This means that it would not be necessary to know the relative proportions of IR and visible light at the time the image is acquired. In this embodiment, stored values for irises may be stored for visible and IR illuminations. Then the desired proportion may be generated using a portion of the stored iris values for visible and IR light. Recognition is initiated at, for example, a 50 to 50 ratio of IR light to visible light. A rough iris match is found. Then the ratio is modified to make the match precise

In Figures 3, 4a, 4b, 5a, and 5b, one placement of the light sources is to have the WFOV light source coaxial with the WFOV imager, another WFOV light source slightly off-axis from the same WFOV imager, a NFOV illuminator that is either a static panel beneath the mirror unit, or co-axial with the NFOV imager and the mirror. The two WFOV illuminators are turned off and on alternately. The alternated images will be very similar, except that red-eye will be very apparent in the coaxial image, and less apparent in the other image. An image alignment and subtraction (or simply subtraction) will yield the eye position. Polarized filters may also be used in the system to enhance the images acquired by the system to aid in locating and identifying the user's facial features in the images. A rotational polarized filter can also be used.

It is possible to have steered illumination that is not coaxial with the NFOV imager. In this embodiment there is a mixture of the static panel NFOV illuminators and a focused co-axial illumination. The light panel could then presumably be smaller and less powerful.

The NFOV imager 14 and WFOV imagers 10 may be used as stereo imagers. For example, a WFOV imager could locate the x, y position of the eyes. The depth would be unknown, but it is known that the user is within a certain range of depths. The NFOV imager could be pointed to the mid-point of this range and a WFOV image

acquired. The mirror could then be rotated precisely by a known angle by the stepper motor and a second WFOV image could be acquired. The two WFOV images could then be aligned and then the location of the NFOV illumination beam that is visible in the WFOV images could be recovered. The displacement of the NFOV beam in the NFOV images can then be used to compute the depth of the user and to drive the mirror 16 to the eye.

In either of the configurations of Figures 4a and 4b or Figures 5a and 5b, the light is directed by the mirror 16 to the area in space which is being imaged by the NFOV imager 14. This is desirable because it allows the heads and eyes of the persons being imaged to be uniformly illuminated regardless of their respective positions in the range which may be imaged by the NFOV imager 14. Once an approximate head position is known, this light source may be used to provide a known level of illumination to the head portions of the images captured by the WFOV imagers 10 and 12. This illumination scheme relies on the X, Y, Z coordinate map that exists between the mirror 16 and the WFOV imagers 10 and 12.

Figures 6 through 13 illustrate functions performed by the WFOV imager 10 and/or NFOV imager 14 when it is being used as a medium field of view imager, responsive to the zoom portion of the signal F/Z. In Figure 6 an image 610 returned by the WFOV imager 10 is a low resolution image of 160 horizontal pixels by 14 vertical pixels. The WFOV image may be captured after the system is alerted that a user is present by inserting a card into the ATM or by other means, such as by a conventional proximity detector or by continually scanning the WFOV imager for head images. It is contemplated that the system may identify a customer with sufficient accuracy to allow transactions to occur without using any identification.

In Figure 6, the image 610 is examined to locate the head and eyes of the user. A method for locating the head illustrated in Figure 6 makes use of flesh tones in the image. The image 610 is scanned for image pixels that contain flesh tones. The image returned by the imager 10 includes a luminance component, Y, and two color difference components, U and V. Whether flesh tones exist at a particular pixel position can be determined by calculating a distance function (e.g. the vector magnitude in color space) between the U and V components of a trial pixel position and a predetermined pair of color component values, U0 and V0, which are defined to represent flesh tones. If the pixel is within a predetermined vector distance of the U0 and V0 values (e.g. if the vector magnitude is less than a threshold), then the pixel position is marked. After marking the image for flesh-tone pixels, the stereo module 316 (shown in Figure 3) defines rectangular regions, 612 and 614 each of which includes at least some percentage, P, of flesh-tone pixels. The values of U0 and V0 vary with the imager 10

and light sources that are used. The percentage P is 60 but may need to be adjusted around this value for a particular implementation.

The regions 612 and 614 are then analyzed by the stereo module 316 to locate relatively bright and dark regions in the image as is illustrated in Figure 6 and described below with reference to Figures 15 and 17. Each region is assigned a possible type based on its relative brightness in the image; relatively dark regions surrounded by relatively bright regions may be defined as potential eyes, while a relatively bright region may be defined as a potential cheek or forehead. Other criteria may be established for recognizing nose and mouth regions. In Figure 7, regions 710, 712 and 714 are classified as a potential forehead or cheek regions and regions 716 and 718 are classified as potential eye regions.

The classified regions are then passed to the stereo module 316, shown in Figure 3. This process uses a symmetry algorithm, described below with reference to Figure 18a, to determine if the classified regions are in relative positions that are appropriate for a face. In the segment 612, for example, a potential forehead region, 710 is above the two potential eye regions, 716 and 718 which, in turn, are above potential cheek regions 712 and 714. In addition, the two potential eye regions 716 and 718 are roughly horizontal. Thus, the segment 612 of the WFOV image 610 is recognized as corresponding to a face. An alternative process of locating the user's head and eyes is described below with reference to Figures 18b through 18i.

After the target eye locations have been identified by the stereo module 316, the image 610 may be monitored for several seconds while the NFOV image is being captured and processed. The image 610 is monitored for motion tracking, as described above, and to check for a change in the image at the identified eye position which corresponds to a blinking motion. It is contemplated that a positive recognition may not be allowed to occur until a blinking motion has been detected. This extra step provides a further check that the target locations are eye locations and ensures that the face being imaged is not a photograph.

If multiple facial images are recognized in the WFOV image, the image which is closest, as determined by either the stereoscopic techniques or by the rangefinder 332, is selected as the image to be searched. Alternatively, the closest user in the WFOV images may be identified prior to localizing the user's face. If one or more facial images is equally close to the WFOV imager 10, then the facial image which is closer to the center of the image is selected.

Once a facial image has been selected, the X, Y coordinates of its eyes are passed to the internal database 318 by the stereo module 316. In addition, the stereo module 316 sends a signal to the control process 310 indicating that an eye location has been determined. If distance (Z coordinate information) is known, either from the

stereoscopic analysis of the images or from the rangefinder 332, it is combined with the X, Y information in the internal database by the process 310.

In response to the signal from stereo module 316, the process 310 signals the NFOV imager 14 to obtain images of the eyes using the stored X, Y coordinate positions. The imager 14 captures two images of the eye, a relatively low resolution image (e.g. 160 by 120 pixels) which may be used by the NFOV/ medium field of view tracking process 322 to locate the eye and to focus the imager 14, as described below with reference to Figure 16b. The imager 14 also obtains a high-resolution image (e.g. 640 by 480 pixels) which is processed by the preprocessor 324 to separate the iris portion of the image for processing by the iris classification and comparison process 326. The steps performed by the processes 322 and 326 are illustrated in Figures 8 through 12.

Although all of the steps outlined above are disclosed as being performed using the WFOV imager, it is contemplated that the steps of capturing the image segment 612 and the processing steps described with reference to Figure 7 may be implemented using the NFOV imager 14, operating as a medium field of view imager under control of the signal F/Z.

Figure 8 shows two eye images, 810 and 812, obtained by the NFOV imager 14 in response to the eye location information provided by the stereo module 316 (shown in Figure 3). Even though the WFOV imager 10 and NFOV imager 14 are calibrated in the X, Y, Z coordinate space, the NFOV imager may not find the user's eye at the designated position. This may occur, for example, if the user moves between when the WFOV and NFOV images are captured or if the Z coordinate information is approximate because it is derived from the sonic rangefinder 332. In these instances, the tracking process 322 may scan the NFOV imager in the X or Y coordinate directions, as indicated by the arrows 814 and 816 in Figure 8, or change its focus to scan the image in the Z coordinate direction, as shown by the arrows 910 and 912 of Figure 9.

The eye may be found in the NFOV image by searching for a specular reflection pattern of, for example, the light sources 126 and 128, as shown in Figure 11. Alternatively, a circle finder algorithm, such as that described below with reference to Figure 19 may be executed. In another alternative embodiment, an autofocus algorithm, described below with reference to Figure 21, which is specially adapted to search for sharp circular edges or sharp textures may be implemented in the tracking process 322 to focus the low-resolution image onto the user's eye.

In an alternative embodiment, the position of the user's eye ZEYE could be modeled using the equation below.

$$ZEYE = q + V$$

q represents the true coordinates of the user's eye and V is a vector of additive noise. If multiple measurements of the location of the user's eye were available over time, the statistical description of V could be used to construct a recursive estimator such as a Kalman filter to track the eye. Alternatively, a minimax confidence set estimation based on statistical decision theory could be used.

Once the user's eye has been located by the NFOV imager, the resolution of the image is changed to 640 pixels by 480 pixels to obtain a high-quality image of the eye. This high-quality image is provided to the iris preprocessor 324, shown in Figure 3. The first step performed by the preprocessor 324 is to rotate the entire image by a predefined amount to compensate for rotational distortion introduced by the mirror 16. In the next step, the specular reflections of the light sources, for example the light sources 126 and 128 are located in the image.

These specular reflections are used to determine the direction in which the user is gazing. If the specular reflections are close to the pupil area, then the user is looking straight at the NFOV imager 14 and the iris will be generally circular. If, however, the specular reflections are displaced from the pupil region then the user is not looking at the imager 14 and the recovered iris may be elliptical in shape. In this instance, an affine transformation operation may need to be applied to the iris to convert the elliptical image into a roughly circular image. The type of operation that is to be applied to the recovered iris may be determined from the relative positions of the specular reflections and the edge of the iris in the image. It is contemplated that a similar correction may be made by analyzing the circularity of the image of the iris. Any recognized non-circular iris may be warped into a corresponding circular iris. Alternatively, the computation for the warping of image may be computed from the expected gaze direction of the user and the recovered X, Y, and Z position of the user's eye.

Before the gaze direction can be determined, however, the iris preprocessor 324 locates the pupil boundary 1210 and the limbic boundary 1212 of the iris, as shown in Figure 12. These boundaries are located using a circle finding algorithm, such as that described below with reference to Figure 19. Once the pupil boundary has been located, the image can be corrected to normalize the gaze direction of the user.

The next step is to find horizontal and near-horizontal edges in the image. These edges, indicated by the reference number 1214 in Figure 12, correspond to the user's eyelid. When the pupil boundary, the limbic boundary and the eyelid boundary have been determined by the preprocessor 324, the portion of the image that is contained within these boundaries (i.e. the iris) is extracted from the NFOV image and is passed to the iris classification and comparison process 326. This step is not needed when the process 326 uses the iris comparison method taught by the above referenced

patents to Flom et al. and Daugman. It may be desirable, however, if an iris recognition system based on subband spatial filtering of the iris is used.

The process 326 compares the iris as found by the preprocessor 324 to the irises stored in the customer database 328. If a match is found, the control process 310 is notified and the process 310 notifies the main control and user interface process 339.

Figure 13 illustrates an alternate method of locating the user's eye using the "red eye" effect. This effect is caused by reflection of a light beam from the retina of the user when the light beam is close to the optical axis of the imager. When the customer's pupils are dilated, this effect is especially pronounced and appears as a reddish glow 1310 for a color image or a bright area for black and white image emanating from the pupils. In an alternative embodiment, this effect may be used to quickly locate the user's eyes in the WFOV image. In this alternative embodiment, the light sources in the ATM provide a relatively low level of illumination. When the user places her card in the card reader 134, the close light source 124 is flashed while a WFOV image is captured. This image is scanned for the reddish color or a bright region characteristic of the "red eye" effect. If the color is detected, its location is passed to the stereo module 316 to determine if the relative brightness, location and size of the identified color areas is consistent with a retinal reflection from two eyes. If the stereo module 316 finds this correlation, then the eye positions have been found and they are passed to the internal database 318.

It is contemplated that the specular reflection of the light sources 128 and 126, as shown in Figure 11 could be used in the same way as the "red eye" effect in order to locate the eyes directly in the WFOV image, without first locating the head. In this alternative embodiment, the entire image is scanned for bright spots in the image representing the specular reflection from the light sources. After the reflections are located, their relative brightness, relative position and position in the WFOV image are tested to determine if they are consistent with expected values for a user's eyes.

Figure 14 is a flow-chart illustrating, at a very high level, exemplary steps performed by the control process 310. The control process is initiated at step 1410 when the customer inserts her card into the card slot 134 of the ATM shown in Figure 1c or, as described above, when a possible customer is detected approaching the ATM. Next, step 1412 is executed to capture an image of the user's head and eyes. Details of an implementation of this step are described below with reference to Figures 15 through 18a. As described below with reference to Figure 15, this step may be abbreviated by directly locating the eyes in the WFOV image using either the "red eye" effect or by directly detecting specular reflections, as described above.

Once the eyes have been located, the control process 310, at step 1414, finds the distance between the eyes and the wide-field of view imagers 10 and 12. As set

forth below, this step may not be needed if step 1412 utilized a range map to find the portion of the image corresponding to the customer's head. If range information is not available, it may be calculated at step 1412 using the two WFOV imagers 10 and 12 as described below with reference to Figure 19. If a single imager is used to generate two
5 photometric stereo images a similar technique may be used. If only a single WFOV imager and a single light source are used, range information may be derived from the sonic rangefinder 332, as described above. Alternatively, the range information may be derived in the process of capturing the NFOV image by using either a conventional autofocus algorithm or an algorithm, such as that described below with reference to
10 Figure 20, that is designed to focus on edges and features that are characteristically found in a human eye.

Alternatively, the range information may be obtained at the same time that the head is located in the image. By this method, two stereoscopic images of the customer, captured by imagers 10 and 12 are analyzed to produce a depth map of the entire field
15 of view of the two imagers. This depth map is then analyzed to identify a head as being close to the ATM and being at the top of an object which is in the foreground. Using this method, range information for each point in the image is determined before head recognition begins. Thus, once the head is found in the WFOV images, its distance from the ATM is known.

The next steps in the process 310 locates the iris (step 1416), using the focused NFOV image and then extract a high-quality image (step 1418) for use by the iris classification and comparison process 326. These steps are described below with
20 reference to Figures 21, 22 and 23. The final step in the process, step 1420, recognizes the customer by comparing her scanned iris pattern to the patterns stored in the customer database 328.

Figure 15 is a flow-chart illustrating details of a process that implements the Locate Head and Eyes step 1412. The first step, step 1510, is to capture the WFOV image using imager 10. The next step, 1512, is to locate the user's head in the image. The purpose of this step is to reduce the size of the image that needs to be processed to
30 locate the eyes. This step is performed as outlined below with reference to Figure 16.

Once extraneous material in the image surrounding the user's head has been removed, the process of locating the user's eyes begins at step 1514 of Figure 15. In this step, the process shown in Figure 15 generates an average luminance map of that part of the image which has been identified as the user's head. This may be done, for
35 example, by averaging each pixel value with the pixels that surround it in a block of three by three pixels to generate an average luminance value for each pixel. This averaged image may then be decimated, for example by 4:1 or 9:1. Once the luminance map has been generated, the next step, step 1516, analyzes this map to identify possible

facial features. Details of a process used by step 1516 are described below with reference to Figure 17. Once the possible facial features have been identified, the next step, 1518, uses symmetry analysis, in a manner described below with reference to Figure 18a, to determine which of the features that were identified as possible eyes are most likely to be the eyes of the user.

If, at step 1520, none of the "possible eye" facial features are identified as being the actual eyes, step 1522 is executed to determine if the process of locating the head and eyes should be retried with a new WFOV image. If so, the control process 310 may prompt the user, at step 1524 via a message on the display 110, to look at a target which will place her eyes in a better position to be imaged. After step 1524, the process repeats with step 1510, described above. The step 1522 may not allow unlimited retries. If for example, a retry limit has been exceeded, step 1522 may pass control to step 1524 which notifies the step 1412 of Figure 14 that the system was unable to locate the user's eyes. In this instance, the control process 310 may abort, allowing the user to access the ATM without verifying her identity or the process may attempt to locate the user's eyes by a different method, such as that illustrated by the steps 1530, 1532 and 1534 of Figure 15.

At step 1530, the control process 310 causes the lighting controls 330 to flash the close light source 124 while concurrently causing the WFOV driver 312 to capture a WFOV image. Next, step 1532 scans the WFOV image for color components and associated luminance components which correspond to the retinal reflections that cause the "red eye" effect. Step 1534 verifies that these are appropriate eye locations using a process (not shown) which determines if the size and relative position of the potential eye locations are appropriate and if their position in the image is consistent with that of a person using the ATM. Although steps 1530, 1532 and 1534 are shown in the context of a detector for the "red eye" effect, it is contemplated that similar steps may be used with a system that directly locates eyes in an image using specular reflection. In the modified algorithm, the shaped light sources 126 and 128 are turned on while the WFOV image is captured. The search step 1532 then searches for specular reflections in the image and compares the located reflections with the shapes and relative positions of the light sources 126 and 128. Step 1534 of the modified algorithm is essentially the same as for the "red eye" detector. It is also contemplated that the light sources 126 and 128 do not have to be shaped and that the specular reflections may be detected based on the amount of energy produced in the image at the positions of specular reflections in the image.

Alternatively, it is contemplated that one or more unshaped flashing light sources could be used to detect specularities in the WFOV image. Sequentially captured images would be compared and only those specular reflections showing the

same temporal light/dark pattern as the flashing light sources would be identified as potential eyes.

Whenever relatively bright light sources or flashing light sources are used, it is contemplated that the light sources may be infrared light sources to minimize the discomfort to the customer. Alternatively, if flashing visible light sources are to be used, it is contemplated that the flash rate may be set relatively high, for example, 60 flashes per second, and the appropriate imager may be synchronized with this flash rate to obtain images both when the light is turned on and when it is turned off.

The detection of specular reflections and the "red eye" effect may be combined with each other and with the location scheme described above with reference to steps 1510 through 1518. One way in which these methods may be combined is to replace steps 1514 through 1518 with steps 1532 and 1534, thus scanning only parts of the image which have been identified as corresponding to the user's head for the "red eye" effect or for specular reflections.

Whichever method is used to locate the eyes, if at step 1520 it is determined that a good candidate eye location has been found, step 1526 is executed to establish X, Y coordinate locations for the eyes in the WFOV image. These coordinates can then be converted into coordinates for the NFOV imager 14 using the coordinate map that was generated during the calibration process, described above. At step 1528, the locate head and eyes process terminates.

The process demonstrated by the flow chart shown in Figure 15 may augmented with a process shown in Figure 24 to locate a person who is to be identified using iris recognition from a number of people in, for example, a line and determine the distance that person is from the system for adjustment of the NFOV imager.

The location of a person with respect to, for example, an ATM is located using the WFOV imagers 10 and 12. The process shown in Figure 24 separates people in line at the ATM and finds the distance of the next person from the ATM in order to focus the lens of the NFOV imager 14. Stereo imaging is used to recover the depth and perform head finding. First, a horopter is located in the image by locating the nearest peak in a correlation surface. Alternatively, multiple horopters could be found to reduce false detections caused by backgrounds that produce a large number of false detections. By locating the nearest peak, the closest person in front of the ATM is selected even if there is a larger person standing behind the next person in line.

The process shown in Figure 24 has two major functions: (1) to find the distance of the user's eye from the system and (2) to extract a region of interest (ROI) of the WFOV images containing the user's head. This ROI is used for eye-finding. The WFOV head-finding process has three major steps. The first step is to find a

suitable horopter in the WFOV image, the second step is to compute a disparity map, and the third step is to locate the user's head.

Producing a disparity map using the WFOV images may be computationally intensive. In order to minimize the amount of search that must be performed, the
5 inventors have determined that it is useful to know the approximate overall disparity of the object which is being mapped. By finding the approximate disparity of the user's face region, the search may be limited for correspondences to a small region around the horopter. Using pyramid processing, the horopter is found at low resolution, while the reduced-search disparity map is constructed at a higher resolution.

10 Potential horopters correspond to sufficiently large regions in which all pixels share similar disparities. As opposed to an image alignment approach which selects the horopters corresponding to the largest such region, a horopter is selected that corresponds to the region of greatest disparity. In other words, the horopter of the closest object to the WFOV imagers 10 and 12 is selected. This horopter generally
15 corresponds to the current user's head. By reducing the disparity-map computation to disparities close to this horopter, the current user is separated from the queue of people waiting behind the user. This process works even if the current user occupies far less image area than someone else in the queue.

Based on the disparity associated with the horopter, a disparity map is produced
20 at higher resolution. From the map, it may be determined which pixels in the image correspond to points at approximately the depth of the horopter. This is precursor to the head-finding step which segments these pixels. The disparity search region consists of small disparities around the nominal disparity shift corresponding to the horopter. In this way, the process accommodates large users for whom the horopter
25 may correspond to some plane through the torso (a few inches in front of the face) rather than through the face. By keeping the search region small, the process effectively separates the user from the background.

The head-finding step accepts the disparity map as input and searches within the disparity map for a region with head-like properties. The corresponding region in
30 image coordinate is passed to a template eye finding process described below with reference to Figures 18a through 18c.

The process of locating an individual shown in Figure 24 begins at step 1310, WFOV imager 12 acquires a first image and, at step 1320, WFOV imager 14 acquires a second image. At steps 1330 and 1340, Laplacian pyramids of the first and second
35 image are respectively produced. At step 1350, coarse Laplacian images of the first image and the second image are shifted with respect to each other by a nominal amount to ensure that the horopter search region corresponds to a 3D volume in which the user's head is expected. The images are bandpass filtered. At step 1360, the shifted

first Laplacian image and the shifted second Laplacian image are multiplied to produce a multiplied image. At step 1370, the multiplied images are blurred and subsampled at a level five Laplacian image using a Gaussian pyramid. At step 1380, it is determined if the coarse Laplacian images have been shifted by X samples. X is ten to fifteen samples. For example, the images are shifted in the x direction in one pixel increments from -7 pixels through +7 pixels. For each shift, a product image is formed at step 1360. The result is a set of fifteen product images, each 80 pixels by 60 pixels.

If the coarse Laplacian images have not been shifted by X samples, than, at step 1390, the coarse Laplacian images of the first image and the second image are shifted with respect to each other by another sample. Step 1360 is repeated. Otherwise, at step 1400, all of the blurred and subsampled images are compared to identify the image with the greatest cross correlation peak. Sums of the pixels in each product image is produced yielding a 15-point 1D sampled correlation function. The sampled function is used to determine the nearest peak having a disparity corresponding to the desired horopter.

A "peak" is a sample point whose value is greater than its two neighbors. The peak should also satisfy the added constraint that the value at the peak is at least 25% greater than the mean value of the correlation function. This threshold is determined heuristically to eliminate peaks of small curvature that are results of noise on the correlation function. Once the appropriate peak has been located, the symmetric triangle interpolation method is used to determine the disparity of the horopter to the sub-pixel level equation (2) below.

$$\text{disparity of horopter} = -7 + i + \frac{f_{i+1} - f_{i-1}}{2(f_i - \min(f_{i+1}, f_{i-1}))} \quad (2)$$

i is the index of the peak in the range of, for example, zero through fourteen, and f_i denotes the ith value of the sampled correlation function. The inventors have found that the symmetric triangle approach is superior to several other interpolation approaches such as quadratic splines in cases where the sampling rate of the correlation function is near critical.

The disparity of the horopter is refined to a Gaussian level 2 resolution: A coarse Gaussian level 3 disparity value is used as a nominal shift of one of the WFOV images to the other WFOV image. Product images are constructed by performing shifts of -1, 0, and 1 pixels at level 2 in both the x and y directions. The centroid of the resultant 9-point sampled correlation surface is used as the level 2 horopter disparity.

At step 1410, the greatest cross correlation peak is selected as the closest object. The result at each shift value is a cross-correlation value that indicates the

similarity of the features of the images at a particular shift. Therefore, if the person is at a distance that results in a disparity of 12 pixels at Level 0, then the cross-correlation will yield a peak at $12/(2^2) = 3$ pixels at Level 2. A sub-pixel interpolator allows sub-pixel disparities to be estimated. If a second user is at a further distance resulting in a disparity of 6 pixels at Level 0, then the cross-correlation will yield a peak at $6/(2^2) = 1.5$ pixels at Level 2. If both users are present, then the system locates the peak corresponding to the nearest object which is the person having a cross correlation of 3.

Next, at step 1415, a higher resolution image is selected around the horopter. Finally, at step 1420, the selected cross correlation maintained to be used to access data in a look-up table and a region of interest containing the user's head is extracted from the WFOV images. Alternatively, the selected cross correlation is converted to a depth value. Once the disparity of the horopter through the user's face has been determined, the approximate distance z of the user's face from the WFOV imagers 10 and 12 may be produced using the equation below for nominally unmerged imagers.

$$\text{disparity} = \frac{a}{z} + b$$

a and b are constants and z is the distance between the user and the imagers.

The region of interest containing the user's head is determined by performing stereoscopic analysis on the WFOV images. Points in the image that exhibit disparities close to that of the horopter are identified to produce a disparity map. The region of the image that is examined to locate the disparities is limited to plus or minus four pixels at level 2 of a Gaussian pyramid with respect to the horopter. Each pixel shift, one pixel on either side of the horopter corresponds to approximately 1 inch of depth. The disparity map computation comprises two steps: (1) correlation and (2) flow estimation.

Correlation is performed in single pixel shifts at level 2 of the Gaussian pyramid. This yields nine product images, each 160 pixels by 120 pixels. Correlation surfaces are computed by integrating over 8 pixel by 8 pixel windows around each pixel. For example, this is performed by computing a Gaussian pyramid of the product images, down to level 5 double-density. This is accomplished by oversampling the resulting image. For example, if the resulting image produced is a 40x30x9 correlation surface, it is oversampled to produce a 80x60x9 surface.

Next, flow estimation is performed by finding the greatest peak in each 9-point sampled correlation function. In addition, a confidence value associated with the peak is produced based on the difference between the peak value and the next-highest non-

neighboring value. The symmetric triangle method described above is used to interpolate the disparity to sub-pixel accuracy for peaks above the confidence value.

The disparity map is used to locate the user's head in the image using the process shown in Figure 24a. At step 1420a, a histogram of disparity values is produced. Each group in the histogram constitutes a disparity of 0.5 pixels. The closest, i.e. the highest disparity, group of pixels containing a minimum number of pixels is used as a fine-tuned estimate of the disparity of the face. This step is useful for cases in which the user has his arms stretched out towards the system or a second person in line is peering over the user's shoulder — i.e. in cases in which more than just the user's face falls within the disparity search region. This step limits the useable disparity range. If the threshold is set to the expected number of pixels on the of the user's face at level 5 of the Gaussian pyramid at double-density, then the user's face may be distinguished from among the clutter.

At step 1420b, a 1D projection of the disparity map onto the x -axis is formed. In other words, a histogram of pixel-count is computed in the x -direction, only for pixels within the fine-tuned disparity range. A horopter-dependent variable is computed as the expected size of the face in the x -direction. Similarly, a horopter-dependent threshold is computed as the expected number of pixels of a face. A window is then used to find candidate x locations for the head within the x -histogram. The total number of pixels in the window is checked to determine if it exceeds the horopter depended threshold.

At step 1420c, for each candidate x location of the head, a similar search for the face is performed in the y -direction. A y -histogram is constructed by projecting onto the y axis only those pixels within the fine-tuned disparity range and within the x -limits defined by the window of step 1420c. In this case, the expected height of the user's face in image coordinates is produced based on the expected number of pixels of a face in the y -direction. In other words, the expected height corresponds to the height of a user's head which may be determined from the height of an average user's head. Blobs of pixels which pass both the x -histogram and y -histogram steps are considered valid faces.

At step 1420d, the centroid c of the blob is found. Multiple iterations of centroid-finding are performed using c as the center and the region within the expected width and height of the user's face. This allows the user's face to be found with high accuracy. At step 1420e, the average disparity of pixels of the centered region is computed. This average is used to compute the z value (distance) of the user's eye. Next, at step 1420f, the ROI is produced by selecting the center c as the center of the user's face and extracting a region surrounding the center c as the ROI. A non-separable blob-detector could be used to detect the user's face in backgrounds that

produce a high number of false matches. At step 1420g, it is determined whether the centroid has been found three times. If not, step 1420d is repeated. Otherwise, at step 1420h, the ROI is provided as an output.

The depth value is used to locate the person's head in a 3-D space and focus the NFOV imager on the user. Figure 25 illustrates the process of locating the user's head. Once the user's head has been located, this data may be used to retrieve adjustment values from a LUT to adjust the NFOV imager. As is shown in Figure 25, at step 3700, the depth of the user's head is found. At step 3710, the height of the person's head and lateral location of the person's head are identified. At step 3720, the depth, height and lateral position of the person's head are used to identify which cube in a 3D space in front of the WFOV imagers contains the person's head. Once the cube has been identified, at step 3730, LUT values are obtained from the LUT which correspond to the identified cube. At step 3740, adjustment values for the depth, height, and lateral position of the person's head are calculated. Then, at step 3750, the NFOV imager is adjusted using the adjustment values.

For example, for a height of one hundred ($x=100$), a lateral position of two hundred ($y = 200$) and a depth of nine hundred ($z = 900$); the exemplary adjustment values generated using retrieved values from the LUT are focus = 1200; zoom = 600; ptu_tilt = 2056; ptu_pan = 100 (ptu is the pan and tilt unit). For $x = 105$; $y = 250$; and $z = 920$; the exemplary adjustment values are focus = 1179; zoom = 605; ptu_tilt = 4134; and ptu_pan = 95. For the values $x = 95$; $y = 305$; and $z = 880$; the exemplary adjustment values are focus = 1220; zoom = 590; ptu_tilt = 6250; and ptu_pan = 107. The expected size of the user's iris may also be stored in and provided by the LUT.

Figure 24b shows an alternative process for the process shown in Figure 24. These processes are the same except for steps 1370, 1400, and 1410 which processes the image at a lower resolution.

Alternatively, a process could be used that measures the distance between specularities that are present on the user's eye to determine the distance the user is from the system. In this embodiment, light sources are provided to project light that would create, for example, horizontal lines, i.e. specularities, that appear on the user's cornea. The distance between the lines on the user's cornea varies in relation to the distance that the user is from the system. Because the distance between the light sources is known and constant, the system may measure the difference between the lines that appear on the user's cornea to determine the distance between the user and system. The system may be calibrated by placing an object at a known distance from the system and measuring the difference between the lines created by the light sources on the object. This data is then stored in a LUT that is used to convert the measured distance between the lines into a depth value.

The values stored in the LUT are captured in a calibration process. In the calibration process, a mannequin or other target is placed in different locations within an 3D grid at points in front of the ATM. In the example below, three points are given. In practice, up to 96 points may be taken during the calibration process. The mannequin or target is positioned at points traversing in the x direction (height) keeping the other parameters approximately constant. The mannequin or target is located at a point x, y, z which is at an approximate 3D grid point. The pan tilt unit is moved until the right (or left) eye of the mannequin is in view. The image is focused and zoomed so that the iris of the mannequin or target is at a desired value. For example, the image is zoomed and focused so that the iris comprises approximately 200 pixels in diameter of the image. This is repeated for all points.

The eight points that are at the corners of each cube in the 3D grid are processed to fit a linear equation to locate x, y, and z with focus as the output. The equation minimized is:

$$\text{error} = ((a + bx + cy + dz + exy + fxz + gyz + hxyz) - \text{focus_measured})^2$$

where (1) a, b, c, d, e, f, g, and h are parameters recovered in the minimization and which are stored in the LUT after they are determined, (2) x, y, and z, are coordinates in the 3-D space, and (3) the focus_measured is the measured focus at the 3-D coordinate. Based on the above equation, the equation (3) below is derived:

$$\text{focus_estimate} = a + bx + cy + dz + exy + fxz + gyz + hxyz \quad (3)$$

The above equation is used in real time by the system to estimate the focus anywhere in the cube. The focus_estimate is the calculated focus for coordinates x, y, and z and a, b, c, d, e, f, g, and h retrieved from the LUT once it has been determined which 3-D cube contains the coordinates x, y, and z.

The imager calibration is repeated for zoom, pan_tilt, pan_zoom, and iris aperture if necessary. This results in an equation for each x, y, and z cube for focus, zoom, pan_tilt, pan_zoom, and iris aperture.

The calibration process is illustrated in Figure 26. At step 3800, the mannequin is positioned at a coordinate (XYZ) in a 3-D grid in front of the imager. At step 3810, the pan tilt unit is adjusted until the right (or left) eye of the mannequin is in view of the imager. At step 3820, the focus and zoom of the imager is adjusted on the eye of the mannequin. At step 3830, the pan tilt unit adjustments and focus and zoom adjustments are stored. In addition, other adjustment values may be stored as required.

At step 3840, it is determined whether the pan tilt unit adjustments and focus and zoom adjustments have been obtained for each (x, y, z) coordinate at the 3-D grid. At step 3870, the mannequin is repositioned to the next coordinate (x, y, z) in the 3-D grid in front of the imager if all of the adjustments have not been obtained for each

point. Otherwise, at step 3850, LUT table values for calculating pan tilt unit adjustments and focus and zoom adjustments with each 3-D cube in the 3-D grid is generated using the stored pan tilt adjustments and focus zoom adjustments using the equations described above. Once these LUT table values have been generated, at step 5 3860, the LUT table values are stored in the LUT. Once calibration is done for a given configuration it may be used in all ATM's having the same configuration.

This process may be automated. There are seven components to be calibrated in the acquisition system. These include WFOV imager point-finding, NFOV imager pan tilt mirror point-finding, autofocus, autozoom, autoaperture, iris-size measurement, and point identification. For example, a calibration chart is placed in front of the system. 10 The WFOV imager point-finding, NFOV imager pan tilt mirror point-finding, autofocus, autozoom, autoaperture, iris-size measurement, and point identification are manually aligned for four points on the chart: top-left, top-right, bottom-left and bottom-right. An auto calibration procedure then aligns the remaining twenty one 15 points on the chart.

For example, the WFOV imagers are calibrated based on the WFOV image locations of the four points. The image coordinates of the remaining twenty one points may be estimated using linear interpolation. This yields a region of interest in which coarse-to-fine positioning may be performed. Standard correlation and flow estimation 20 may be performed relative to an artificial reference image comprised of an array of small black disks on a white background. The input from one of the WFOV imagers is replaced by an artificial image. The output from the other WFOV imager is the position of the black disks to sub-pixel resolution.

The pan-tilt mirror is calibrated using an exhaustive search: The pan tilt mirror 25 panned and tilted in small steps such that the NFOV imager tiles a region large enough to guarantee capture of the point concerned, but not so large as to include any other points. Alternatively, the NFOV imager is zoomed out to form a MFOV image large enough to guarantee capture of the point concerned, but not so large as to include any other points. For either approach, it may prove to be difficult to guarantee that only one 30 point is captured on the chart. To ensure that the proper point is captured, each point be individually bar-coded for identification. For example, a spiral bar-code within the black disk may be used a bar code. In this case, the system may be used to differentiate among points.

Once a point has been captured, a spoke-detector may be used to acquire and 35 center the disk within the NFOV. A coarse autofocus process may be added before using a spoke detection process. Autofocus, centering, and re-zooming may be interleaved to adjust the NFOV imager.

Figure 16 is a flow-chart illustrating a process for implementing the Locate Head step 1512 of Figure 15. The process, via the dashed-line arrows 1628 and 1630, illustrates multiple alternative processes for locating the user's head. In the first step 1610 of the process, the stereo module 316 generates a motion profile for the WFOV image. The motion profile is generated from two successive WFOV images. A difference image is derived by subtracting the second image from the first image. The individual pixel values are then squared and the stereo module 316 finds subsets of the image, which are defined by rectangles that surround any large groups of changed pixel values. The coordinates that define these rectangular subsets are used to obtain possible head images from the WFOV image. Instead of using two successive images obtained by the WFOV imager, it is contemplated that the system may generate the difference image by subtracting the WFOV image obtained at step 1510 (shown in Figure 15) from the WFOV imager from a previously obtained image which was captured when it was known that no person was in the field of view of the WFOV imager 10.

As an alternative to step 1610, step 1611 may be executed to locate a person in the WFOV image. Stereoscopic images from the two WFOV imagers 10 and 12 are used to generate a range map for the scene that is within the field of view of the imagers 10 and 12. To generate this map, corresponding pixels in each of the two images are examined and depth information is assigned to each object in the image based on the respective pixel positions of the object in the two images. Once the depth map is generated, it is relatively easy to locate the customer by her body shape and relatively close position to the imagers 10 and 12. Her head is at the upper portion of her body.

If, at step 1612, the stereo module 316 determines from the size and shape of the motion image or of a body image found in the range map, and from the location of the found image in the WFOV image, that one of the potential head images has a high probability of being the user's head, control is transferred from step 1612 to step 1616. Otherwise, control is transferred to step 1614. Step 1614 scans each of the target head images for flesh-tone pixels. At step 1616, the groupings of flesh-tone pixels are compared. The portion of the image corresponding to the grouping that most closely corresponds in its shape, size and position in the image to a nominal head profile is identified as the user's head.

The next step in the process, step 1618, analyzes the WFOV image at the area selected as the user's head to determine various physical characteristics, such as height, head shape, complexion and hair color. These characteristics can be readily determined given the location of the head in the image and approximate range information determined by the range map from step 1611 or the sonic rangefinder 332. At step 1620, these characteristics are compared to stored characteristics for the user who

recently placed her card into the ATM. If several of the determined characteristics do not match the stored characteristics, the captured image, at step 1624, is stored in the internal database as that of a possible unauthorized user. Because these characteristics can change, at least at this stage, the identification process cannot be used to prevent access to the ATM. Accordingly, after step 1622 or after step 1624, the Locate Head process is complete at step 1626.

As shown by the dashed-line arrow 1628, the generation and analysis of the motion profile may be eliminated from the Locate Head process. In addition, as shown by the dashed-line arrow 1630, the security analysis steps 1618 through 1624 may be eliminated.

Figure 17 is a flow-chart illustrating a possible implementation of the step 1516 of Figure 15 in which possible facial features are identified in the head portion of the WFOV image. As described above, the input image to this process is one in which the luminance component of each pixel value is replaced by an average of the pixel and its eight neighboring pixels.

In the first step 1710 of this process, the stereo module 316 selects a first pixel from the luminance averaged image as a target pixel. Next, at step 1712, the target averaged pixel value is compared with its eight surrounding averaged pixel values. At step 1714, if the luminance level of the target pixel is less than that of a number N of its surrounding pixel values then, at step 1716, the target pixel is marked as being a possible eye or mouth. If, at step 1714, the luminance level of the target pixel is not less than the luminance levels of N of its surrounding pixels, then control passes to step 1718 which determines if the luminance level of the target pixel is greater than or equal to the luminance levels of M of its surrounding pixels. If this test is met then, at step 1720, the pixel is marked as being a possible cheek, forehead or nose. In the exemplary embodiment, N is six and M is seven. After step 1716, if the condition is false, step 1722 is executed which determines if all of the pixels in the image have been processed for feature extraction. If more pixels remain to be processed, control is transferred to step 1710 to select a next target pixel. If no more pixels remain, then control transfers to step 1724 and the Identify Possible Facial Features process is complete.

Figure 18a is a flow-chart illustrating an implementation of step 1518. Locate Eyes by Symmetry Analysis. This process is part of the stereo module 316 of Figure 3. The first step in this process, step 1810, identifies, as a target eye pixel position, the first possible-eye pixel position as determined by the feature identification process, described above with reference to Figure 17. Step 1810 also scans the image in a horizontal direction to determine if the image contains another possible eye pixel, displaced from the target pixel by P pixel positions, \pm Q pixel positions. Values for P

and Q are determined by the distance between the customer and the imager and by the decimation of the averaged WFOV image. If a second possible eye pixel is not found within the search area, control returns to step 1810 to choose another possible eye pixel position as the target pixel position.

5 If, however, a second eye pixel position is found at step 1812, control is transferred to step 1814 and the portion of the image representing points higher on the face than the eye pixel positions is searched for possible forehead pixel positions. Next, at step 1816, the portion of the image representing points lower on the face than the eye positions is searched to locate possible cheek, nose and mouth pixel positions.
10 Whatever corresponding features are found by the steps 1814 and 1816 are compared to a generic template in step 1818. This template defines a nominal location for each of the facial features and an area of uncertainty around each nominal location. If some subset of the identified features is found to fit the generic template at step 1820, the target eye location and its corresponding second eye position are identified as
15 corresponding to the user's eyes at step 1822.

Figures 18b and 18c illustrate an alternative template process for locating the user's head and eye. The template process may be used to implement step 1412 shown in Figure 14 or as a replacement for steps 1514, 1516, and 1518 shown in Figure 15. The template process locates the coordinates of the user's eyes from the ROI containing
20 an image of the user's head. The template process uses a template-based approach and information from various filter kernels. The templates are designed to be scaleable to allow eye-finding for varying head sizes and varying distances from the imager.

Each of the templates is scaled in proportion to the size of the face region being processed. This is accomplished using a disparity process that provides a disparity measure from which the approximate distance of the person from the imager may be calculated. Alternatively, the disparity may be used to access a database to provide this
25 information without converting the disparity to a depth value. The distance between the user and the WFOV imagers 10 and 12 (shown in Figure 1c) is used to calculate a disparity value which is subsequently used to produce a scaling factor to scale the
30 template based on the user's distance from the acquisition system.

The template process has two processing paths (1) an eye-finder process for when the user is not wearing glasses and (2) an eye-finder process for when the user is wearing glasses. The template process is divided into two processing paths because the response caused by the user's glasses often interferes with the detection of the user's
35 facial features. The template process first attempts to locate the user's face assuming the user is wearing glasses. If no face is located, then the template process attempts to locate the user's face assuming the user is not wearing glasses.

Once the various filter outputs are available, the processes 2075 through 2092 and 2010 through 2055 are performed in order for templates that are moved around over the entire image pixel by pixel and all positions passing all processes (tests) are recorded. As soon as one position fails one procedure the remaining procedures are not performed at the same position. The procedures are in roughly increasing order of computational complexity and decreasing order of power of selectivity. As a result, the overall process for localizing the user's eyes may be performed with less computation time. Although in the current embodiment all of the processes need to be passed in order for a face to be detected, a subset of the processes may be performed. It is also contemplated that all or a subset of the procedures may be performed and that the pass/fail response of each of the processes will be considered as a whole to determine if a face has been detected.

In Figures 18b and 18c, the images from the WFOV imagers 10 and 12 are filtered at step 2000 to generate filtered images. The filtered images are derived from an image I that is a reduced-resolution version of the image from the WFOV imagers 116 and 118, obtained via blurring and subsampling using an image pyramid process as described in U.S. Patent No. 5,539,674. This patent is incorporated herein by reference for its teachings on pyramid processing. This process decreases the spatial frequencies contained in the filtered images, and improves the computational tractability of the filtering process. In this embodiment, the WFOV images are reduced by a factor of four in both the X and Y dimensions.

The filtered images include (1) a Laplacian filtered image L, a second derivative in both the x and y directions, of the WFOV images (2) an X-Laplacian image L_x, a second derivative in the x-direction, of the WFOV images; (3) a Y-Laplacian image L_y, a second derivative in the y-direction, of the WFOV images; (4) first derivative images D_x and D_y in both the x and y directions of the WFOV images; (5) orientation maps ORIEN of the WFOV images; and (6) thresholded squared Y-Laplacian images (Ty₂) of the WFOV images.

The Gaussian filter coefficients to be used in the equations below are set forth in relation (4) below. In the notation below G(x) refers G(.) oriented in the x, (horizontal) direction, and G(y) refers to G(.) oriented in the y (vertical) direction.

$$G(.) = [1/16, 2/8, 3/8, 2/8, 1/16] \quad (4)$$

The Laplacian filtered images L are constructed by taking the difference between image I and the corresponding Gaussian filtered version of I according to equation (5) below.

$$L = I - I * G(x) * G(y) \quad (5)$$

The X-Laplacian filtered images L_x are produced by taking the difference between the images I and the Gaussian filtered image in the x-direction $G(x)$ according to equation (6) below.

$$L_x = I - I * G(x) \quad (6)$$

The Y-Laplacian filtered images L_y are produced by taking the difference between the image I and the Gaussian filtered image in the y-direction $G(y)$ according to equation (7) below.

$$L_y = I - I * G(y) \quad (7)$$

The first filtered derivative images in the x-direction D_x and the y-direction D_y are each produced using a 5-tap filter having the coefficients set forth below.

$$D(.) = [1/8, 1/4, 0, -1/4, -1/8]$$

For the x-direction first derivative images D_x , the images I are filtered in the x-direction using the 5-tap filter defined above to produce $D(x)$ and filtered in the y-direction using the Gaussian filter coefficients in relation (4) to produce $G(y)$. The x-direction first derivative image D_x is produced according to equation (8) below

$$D_x = I * D(x) * G(y) \quad (8)$$

For the y-direction first derivative images D_y , the WFOV images are filtered in the y-direction using the 5-tap filter defined above to produce $D(y)$ and filtered in the x-direction using the Gaussian filter coefficients of relation (4) to produce $G(x)$. The y-direction first derivative images are produced according to equation (9) below.

$$D_y = I * D(y) * G(x) \quad (9)$$

The orientation maps ORIEN are computed using equation (10).

$$\text{ORIEN} = \arctan (2 * D_x * D_y / (D_x * D_x - D_y * D_y)) \quad (10)$$

Equation (10) is computed in two steps. First, two image maps $j1map$ and $j2map$ are produced according to equations (11) below.

$$j1map(.) = 2 D_x(.) D_y(.) \quad (11)$$

$$j2map(.) = D_x(.) D_x(.) - D_y(.) D_y(.)$$

Then the orientation map is produced according to equation (12) below.

$$\text{ORIEN} = \text{atan2}(j1map(.), j2map(.))/2 \quad (12)$$

The thresholded squared Y-Laplacian images $Ty2$ of the WFOV images are obtained by squaring the Y-Laplacian filtered images L_y and thresholding all values less than, for example, sixteen. In others words, all values below the threshold are changed to, for example, zero.

In Figures 18b and 18c, the WFOV glasses specularities detection process determines if the user is wearing glasses. As discussed above, the presence of glasses can cause false detection of the user's facial features. For example, the frame of the

user's glasses can provide more energy to the check brightness process 2020, described below, causing process 2020 to produce an erroneous result.

The WFOV glasses specularities detection process 2077 detects and localizes the position of the specularities from the user's glasses caused by the reflection of light incident from the light sources. The glass specularities are detected to determine if the user is wearing glasses, estimate the location of the user's eyes using the detected specularities, and determine what portion of the user's eyes may be occluded by the specularities.

The WFOV glasses specularities detection process 2077 is explained with reference to Figures 18d and 18e. First, at step 2077a, shown in Figure 18e, the right or top eye template is processed to detect specularities. The top eye template and the other templates are shown in Figure 18d.

The templates used in these processes include a number of templates oriented with various regions of a person's face. The templates include (1) left and the right eye templates 720 and 722, (2) nose template 724, (3) bridge of the nose, mid-eye, template 726, (4) left and right cheek templates 728 and 730, and (5) mouth template 732. The bridge-of the nose or mid-eye template 726 is a small template lying in the middle of the user's two eyes. The WFOV imagers 10 and 12 (shown in Figure 1c) are horizontally oriented resulting in facial images being rotated by ninety (90) degrees. The template is designed for this orientation of the user's face. If the image of the user's face is not rotated, the templates are rotated by ninety degrees to compensate for the different orientation of the user's face produced by the WFOV imagers 10 and 12.

The various templates and corresponding nomenclature for each template is provided in the table below.

User's Feature	Reference Description
right eye template 720	top eye
left eye template 722	bottom eye
right cheek template 728	top cheek
left cheek template 730	bottom cheek
nose template 724	nose
mouth template 732	mouth
bridge-of the nose template 726	mid-eye

The sizes of the templates were chosen by warping a set of sample images from ten people on top of each other so that the eyes were all aligned and then measuring the separation of the two eyes and other features of the face. At a mean distance of about

two feet from the WFOV imagers 10 and 12 (shown in Figure 1c) the eye separation of an average user was found to be about seventeen pels in the subsampled imagery used by the template matching process. This assumes an operation range of the acquisition system of one to three feet from the WFOV imagers 10 and 12.

There are two different sets of templates: one for when the user is wearing glasses (the glass template set) and one for when the user is not wearing glasses (the no glass template set). The templates in these sets are not the same because the process for detecting the user's eyes when the user is wearing glasses is different from the process for detecting the user's eyes when the user is not wearing glasses. As a result, not every template is needed for both sets of procedures.

Further, the top and bottom eye templates 720 and 722 are larger in the glasses template than the no glasses template because two specularities can appear on each lens of the user's glasses which may be larger than the user's eye. As is described above, two specularities may be detected because two light sources 126 and 128 may be used to illuminate the user's face. In order to detect these specularities, the size of the eye templates 720 and 722 is increased.

The template definition used when the user is wearing glasses is defined in the tables below.

GLASSES TEMPLATE SET SIZES
(WHEN THE USER IS WEARING GLASSES)

TEMPLATE	TEMPLATE SIZE (in pixels and rows by columns)
right eye template 720	12 by 3
left eye template 722	12 by 3
mouth template 732	12 by 10
nose template 724	6 by 8

GLASSES TEMPLATE SET OFFSETS
RELATIVE TO THE SELECTED PIXEL
(WHEN THE USER IS WEARING GLASSES)

TEMPLATE	OFFSE T (in pixels)
left eye X offset	2
left eye Y offset	22
right eye X offset	2
right eye Y offset	0
mouth X offset	18
mouth Y offset	11
nose X offset	5
nose Y offset	14

5

The template definition for the user without glasses is in the tables below.

NO GLASSES TEMPLATE SET SIZES
(WHEN THE USER IS NOT WEARING GLASSES)

TEMPLATE	TEMPLATE SIZE (in pixels and rows by columns)
right eye template 720	6 by 3
left eye template 722	5 by 3
right cheek template 728	6 by 8
left cheek template 730	5 by 8
nose template 724	6 by 8
mouth template 732	12 by 10
bridge-of the nose template 726	6 by 3

10

NO GLASSES TEMPLATE SET OFFSETS
RELATIVE TO THE SELECTED PIXEL
(WHEN THE USER IS NOT WEARING GLASSES)

TEMPLATE	OFFSET (in pixels)
left eye Template X offset	2
left Eye Template Y offset	17
right Eye Template X offset	2
right Eye Template Y offset	0
nose Template X offset	5
nose Template Y offset	9
left Cheek Template X Offset	8
left Cheek Template Y offset	17
right Cheek Template X offset	8
right Cheek Template Y offset	0
mouth Template X offset	18
mouth Template Y offset	5
mid Eye Template X offset	2
mid Eye Template Y offset	9

5

In Figure 18e, at step 2077a, each of the pixels located in the top eye template 720 is processed to locate two by two pixel regions, four pixels, which have large Laplacian values. For example, values above a threshold value of twenty are determined to be large. If no specularities are detected in the top eye template 720, the next pixel is selected at step 2075 and step 2077 is repeated if it is determined at step 2076 that all the pixels have not already been tested. Otherwise, at step 2077b, the

10

bottom eye template 722 is processed to detect specularities. If specularities are not detected in the bottom eye template, then the next pixel is selected at step process 2075 and step 2077 is repeated. By detecting a specularity in each lens of the user's glasses, false detection of specularities caused by the tip of the user's nose may be prevented.

5 At step 2077c, if two specularities are detected in the right and left eye templates 720 and 722, then the mean position of the two specularities in each eye template 720 and 722 is designated as the location of the user's eye. This position is designated as the location of the user's eyes because the positioning of the two light sources 126 and 128 (shown in Figure 1c) causes the specularities to occur on either side of the user's
10 eye in the eye glass. If one specularity is detected, then its position is designated as the estimated location of the user's eye.

 In figure 18b, the next process, the eye blob process 2080, detects the position of the user's eye. Because a portion of the user's eye may be occluded by the specularities from the user's glasses, the eye-blob process 2080 examines portions of
15 the image of the user's eye which are not occluded. In other words, positions other than the location of the detected specularities are examined. The eye blob process 2080 examines a ten by ten pixel region of the Laplacian filtered image L surrounding the location of the specularities detected by process 2077. A two-by-two pixel template is used to examine the ten-by-ten region of the Laplacian filtered image L to locate
20 negative Laplacian values, for example, values less than negative one. Values less than negative one are considered to be a location corresponding to the user's eye. The location of the Laplacian values smaller than negative one in each eye template 720 and 722 are clustered and averaged to produce an estimated location for each of the user's eyes.

25 The response of the user's eye glass frames in the Laplacian filtered image L is similar to that of the user's eyes. Further, the distance of the eye glass frame from the detected specularities is similar to the distance of the user's eyes from the specularities. Thus, to overcome difficulties that may be caused by detection of the frames, the detected portions of the frames and the detected portions of the eye are used in
30 combination to determine the average position of the user's eye. The eye glass frames may be used in combination because the frames are generally symmetrical about the users eyes and, as a result, do not affect the average position of the user's eye.

 If the locations of the user's eye can not be detected in each of the templates 720 and 722 at step 2080, step 2075 is repeated. Otherwise, the mouth process 2085 is
35 implemented.

 The mouth process 2085 uses the expected orientation of the user's mouth to determine if the user's mouth is detected in the mouth template 732. Thus, false positive detections caused by, for example, the side of the user's head are detected.

5 The mouth process 2085 uses a movable template because the position of a person's mouth varies widely with respect to the user's eyes and nose from person to person. The mouth template 732 (shown in Figure 18d) may be moved in a direction towards or away from the position of the user's eyes to account for this variability. Although the mouth template 732 may move, the mouth template should not be moved to overlap or go beyond the pixel border columns of the image. The pixel border columns are the pixels located at the edge of image. Thus, for example, the mouth template should not be moved to overlap the two pixel border columns located at the edge of the image. The pixel border columns are avoided because of the border effects caused by the size of the filter kernels.

10 The mouth process 2085 calculates the average orientation of pixel values in the mouth template 732 by averaging the j1map and the j2map images over the mouth template region and using the resulting average j1map and j2map to compute the orientation using equation (12). The mouth process 2085 then determines if the average position is no greater than, for example, ten degrees on either side of a horizontal reference line. The horizontal reference line is with reference to an upright orientation of the user's face. In the ninety degree rotated version of the user's head acquired by the WFOV imagers 10 and 12 (shown in Figure 1c), the horizontal line corresponds to a vertical reference line. If the average position is greater than ten degrees, the mouth template is moved and the test is repeated. The mouth template 732 is moved until the user's mouth is detected or each of the possible positions of the mouth template 732 is processed to detect the user's mouth. If the user's mouth is not detected, then process 2075 is repeated. Otherwise, the nose process 2090 is implemented.

25 The nose process 2090 is described with reference to Figure 18f. Steps 2090a through 2090f impose a minimum energy constraint before further processing in steps 2090g through 2090l. At step 2090a, a thresholded squared Y-Laplacian image Ty2 is produced and averaged. The thresholded squared Y-Laplacian image Ty2 is produced by squaring each of the values in the nose template 724 and eliminating the squared values below sixteen. Then the remaining squared thresholded Y-Laplacian values are averaged over the extent of the template. At step 2090b, it is determined whether the averaged value is greater than a threshold of, for example, 45. This process ensures that a minimum number of pixels in the Y-Laplacian image Ly have a magnitude greater than four. Pixels having values below four are not considered. At step 2090c, the nose template 724 is moved if the average value is less than the threshold. The nose template 724 is movable in both the horizontal and vertical directions. At step 2090c the nose template is allowed to move four pixels vertically on either side of the image now halfway between the two eyes, the first direction.

At step 2090d, the nose template position is stored if the average value is greater than the threshold. At step 2090e, if each of the positions of the nose template 724 has been processed, then processing passes to step 2090f. Otherwise, step 2090c is repeated. At step 2090f, if none of the nose templates 724 positions is stored and do not satisfy the minimum energy constraint, then processing passes to process 2075, shown in Figure 18b.

Otherwise, at step 2090g, for each nose template 724 position that satisfies the initial nose energy criterion, the average Y-Laplacian energy and the average X-Laplacian energy is computed. In this implementation, this is done by averaging $L_y \times L_y$ and $L_x \times L_x$, respectively, over the extent of the template. Other energy measures such as absolute value could also be used. At step 2090g, if the Y-Laplacian energy is greater than a threshold, then step 2090h is implemented, otherwise step 2090j is implemented. At step 2090h, each of the average Y-Laplacian energies is compared to a corresponding average X-Laplacian energies for each stored template position. At step 2090k, the nose template position having the largest ratio of the average Y-Laplacian energy to the average X-Laplacian energy is selected. Steps 2090g through 2090i are repeated for each of the nose template 724 positions along a second direction moving horizontally towards and away from the user's eyes at positions corresponding to the stored nose templates 724.

At step 2090i, it is determined for the largest ratio whether the average Y-Laplacian energy is greater by a factor of Z than the average X-Laplacian energy. The factor Z is, for example, 1.2. The user's nose is expected to have a higher Y-Laplacian filtered image L_y response than an X-Laplacian filtered image L_x response because of the vertical orientation of the user's nose. In the oriented image, the x-direction extends from the user's chin through his nose to the top of his head.

In Figures 18b and 18c, if none of the templates has a ratio greater than the factor Z, then processing passes to process 2075. Otherwise, at step 2090l, the template position having the largest ratio is selected as the position of the user's nose.

At step 2092, after each process 2075 through 2090, the positions of the templates are stored in a data structure keeping track of all template positions obtained. Steps 2075 through 2090 are repeated for different template positions. As a result, there may be more than one detected position for the user's eyes. At step 2005, after all of the pixels have been tested, it is determined if any eye positions have been stored at step 2092. If so, at step 2060, the detected positions are clustered to more accurately localize the position of the user's eyes.

The detected positions are clustered according to the following process. The positions within a certain radius are considered to belong to the same cluster. First, a center of the cluster is selected and each detected position is added to the cluster within

a specified radius of the cluster. The specified radius is, for example, one fourth of the separation of the extreme ends of the two eye templates 720 and 722 (shown in Figure 18d). New clusters are seeded and the process is repeated. The number of positions located within each cluster is stored in a database.

5 At step 2065, the position having the best coordinates is determined and stored for each cluster. The best coordinates are the coordinates that have x-coordinate closest to the user's mouth. In other words, the lowest point in an upright image. By using this measure, bad localization of the center of the user's eye caused by a large number of detections on the user's eyebrows when the eyebrows are close to the user's eyes is reduced. The y-coordinate is set to be the same as the center of the cluster. Finally, at 10 step 2070, the location of the user's eye is selected as the "best coordinates" of the cluster that has the greatest number of members in the cluster. Alternatively, more than one, for example, three of the clusters having the greatest number of members may be stored for further processing.

15 Next, at step 2071, the WFOV corneal specularities detection process is implemented. Process 2071 refines the position estimated by the template-based eye finding algorithm. Such a refinement is important because the template-based algorithm, as described above, extracts eye coordinates in a reduced-resolution version of the WFOV image. This reduction in resolution means that the eye is not localized in 20 the WFOV image as precisely as possible. Therefore, the WFOV corneal specularities detection process detects specularities on the cornea in the full-resolution WFOV image and uses these to refine the position of the template-extracted eye. Process 2071 is described below with regard to Figure 18j. At step D1000, a pixel CS1 is selected from the WFOV image. Next, at step D1010, the corneal specularity CS1 coarse 25 detection process is implemented. Process D1010 is described below with reference to Figure 18k.

 Process D1010 determines whether pixel CS1 is likely to be a small specularity of the user's eye. It does so by testing whether the pixel is brighter than its neighbors on three of its sides. Specifically, at step D1010a, it is determined whether the pixel 30 two units below CS1 is at least darker than the selected pixel CS1 by a threshold. If not, processing passes to step D1000. Otherwise, at step D1010b, it is determined if the pixel two units to the left of the selected pixel CS1 is darker than the selected pixel by a threshold. If not, processing passes to step D1000. Otherwise, at step D1010c, it is determined if the pixel two units to the right of the selected pixel CS1 is darker than 35 the selected pixel by a threshold. If so, processing passes to step D1020 shown in Figure 18j.

 In the preferred implementation, there are multiple bright illumination sources. Each of these sources produce a specularity on the user's cornea. Therefore, the

corneal specularities rather than just one to reduce the number of false detections. Furthermore, because of the known approximate orientation of the user's face in the WFOV imagery, and the knowledge of the position of the illuminators, the expected geometric configuration of the specularities is known and can be used to make the WFOV corneal specularities detector more selective. In the embodiment, there are two illuminators spaced horizontally, Therefore, the WFOV corneal specularities detector, after finding one potential specularities CS1, attempts to find a second specularities positioned below it.

Specifically, at step D1020, a cone shaped search region is selected with respect to pixel CS1 for searching for another pixel CS2 that corresponds to a second specularities in the WFOV image. The cone shaped region CONE is shown in Figure 18l. Next, at step D1030, the corneal specularities CS2 coarse detection process is implemented in the cone search region as described below with regard to Figure 18m. At step D1030a, it is determined if there are any pixels in the cone shaped region remain to be processed. If not, step D1000 (shown in Figure 18j) is repeated. Otherwise, at step D1030b, it is determined whether the pixel two units above CS2 is at least darker than the selected pixel CS2 by a threshold. If not, processing passes to step D1030a. Otherwise, at step D1030c, it is determined if the pixel two units to the left of the selected pixel CS2 is darker than the selected pixel CS2 by a threshold. If not, processing passes to step D1030a. Otherwise, at step D1030d, it is determined if the pixel two units to the right of the selected pixel CS2 is darker than the selected pixel CS2 by a threshold. If so, processing passes to step D1040 shown in Figure 18j. Otherwise, step D1030a is repeated.

Processes D1040, D1050, and D1060 further verify that pixels CS1 and CS2 are corneal specularities, i.e. bright points rather than bright lines. Process D1040 is described below with reference to Figure 18n. At step D1040a, it is determined whether all pixels on a line two units to the left of and up to two units below the pixel CS1 from the coarse detection process D1010 (shown in Figure 18j) are darker than pixel CS1 by at least the threshold. The line is shown in Figure 18o. If not, step D1000 (shown in Figure 18j) is implemented. Otherwise, at step D1040b it is determined whether all pixels on a line two units to the right of and up to two units below the pixel CS1 are darker than pixel CS1 by at least the threshold. The line is shown in Figure 18p. If not, step D1000 (shown in Figure 18j) is implemented. Otherwise, at step D1040c it is determined whether all pixels on a line two units two units below and up to two unit on either side of the pixel CS1 are darker than pixel CS1 by at least the threshold. The line is shown in Figure 18q. If not, step D1000 (shown in Figure 18j) is implemented. Otherwise, the corneal specularities process for pixel CS2 D1050 is implemented.

Process D1050 is described below with reference to Figure 18r. At step D1050a, it is determined whether all pixels on a line two units to the left of and up to two units above the pixel CS2 from the coarse detection process D1020 (shown in Figure 18j) are darker than pixel CS2 by at least the threshold. The line is shown in Figure 18s. If not, step D1000 (shown in Figure 18j) is implemented. Otherwise, at step D1050b it is determined whether all pixels on a line two units to the right of and up to two units above the pixel CS2 are darker than pixel CS2 by at least the threshold. The line is shown in Figure 18t. If not, step D1000 (shown in Figure 18j) is implemented. Otherwise, at step D1050c it is determined whether all pixels on a line two units two units above and up to two unit on either side of the pixel CS2 are darker than pixel CS2 by at least the threshold. The line is shown in Figure 18u. If not, step D1000 (shown in Figure 18j) is implemented. Otherwise, the corner process D1060 is implemented.

Many unwanted specularities come from the rims of glasses. In general these specularities lie on image structures that are linear or edge-like in a direction over the local image region. Specularities off the cornea however are usually isolated and form a corner type structure.

Corner process D1060 differentiates edges from corners. Consider the Matrix:

$$\begin{matrix} I_{xx} & I_{xy} \\ I_{xy} & I_{yy} \end{matrix}$$

where I_{xx} is the sum of $I_x * I_x$ over a five by five region, I_{xy} is the sum of $I_x * I_y$ over a five by five region, and where I_{yy} is the sum of $I_y * I_y$ over a five by five region where the I_x and I_y are the image gradients in the horizontal and vertical edges computed by convolving with a (-1, 0, 1) filter. The determinant of this matrix is zero in two cases: (1) when there is no structure and (2) when there is linear image structure. When there is corner like structure then the determinant is large. The determinant is:

$$(I_{xx} * I_{yy} - I_{xy} * I_{xy}).$$

The corner process D1060 uses these principles and performs a corner detector process at all positions where specularities are detected by calculating determinants. Process D1060 rejects those specularities where the determinant is less than a threshold which suggests that the specularities lie on the rims of the glasses and not in the cornea.

Finally, at step D1070, location of the corneal specularities is selected as the positions of pixels CS1 and CS2.

In Figure 18c, at step 2072, if it is determined that the positions of the corneal specularities are near the "best coordinates" from step 2070, then at step 2073, the corneal specularities closest to the "best coordinates" are selected as the localization, i.e. positions, of the user's eyes. Thus, the "best coordinates" obtained from the template

matching process that uses low resolution WFOV imagery, are refined by the WFOV corneal specular detector that uses high resolution WFOV imagery. Otherwise, at step 2074, the "best coordinates" are selected as the localization, i.e. positions, of the user's eyes.

5 In Figure 18b, if it is determined at step 2005 that no eye positions are stored then processes 2010 through 2055 are implemented. Process 2010 selects a starting position in the image. Each time process 2010 is repeated, a new pixel in the image is selected to determine the location of the templates within the image. The eye energy process 2015 detects the user's eyes in the top and lower eye templates 720 and 722.
10 The eye energy process is described below with reference to Figure 18g.

At step 2015a, the average Laplacian value of the Laplacian image L in the top eye template is calculated. At step 2015b, it is determined whether the average Laplacian value is less than negative one which indicates the presence of the user's eye in the top eye template 720. The user's eye regions are expected to yield a negative
15 Laplacian response because the eye regions are isolated dark blobs. Other spurious noise responses which could be detected are eliminated by averaging the Laplacian values over each of the eye templates 720 and 722. If the user's eye is not detected in the top eye template, another position is selected at step 2010 (shown in Figure 18b).

In Figure 18g, at step 2015c, once the user's eye has been detected in the top
20 eye template 722, the average Laplacian value of the Laplacian image L of the lower eye template is calculated. At step 2015d, it is determined whether the average Laplacian value is less than negative one which indicates the presence of the user's eye in the lower eye template 722. If the user's eye is not detected in the lower eye template, another position is selected at step 2010 (shown in Figure 18b).

25 In Figure 18g, the lower eye template 722 is adjustable to allow a degree of freedom on either side of the position of the top eye template 720. This accounts for slight variations in the natural poses of the user's head. As a result, more than one possible position of the user's left eye may be detected. Step 2015e selects the lower eye templates 722 having an average Laplacian value less than negative 1. At step
30 2015f, if there are two or more selected lower eye templates 722, then the lower eye template that is positioned closest to a point directly below the position of the top eye template 720 is selected as the position of the user's left eye. The positions of the user's right and left eyes are stored in a database if they are detected in the upper and lower eye templates 720 and 722.

35 In Figure 18b, if the user's eyes are detected, then the cheek brightness process 2020 is performed. Because the user's cheeks are usually bright textureless regions, the cheek brightness process 2020 determines whether the cheek templates 728 and 730 have a minimum average brightness value. The average gray level value is calculated

for each cheek template and compared to a brightness threshold. The brightness threshold for the upper cheek template 728 is, for example, forty five and the brightness threshold for the lower cheek template is, for example, 30. The average brightness value for the lower and the upper cheek templates 728 and 730 is different because the position of the rotated WFOV imagers 10 and 12 is closer to the user's right eye in the right eye or upper eye template 720 than to user's left eye in the left eye or lower eye template 722. Thus, the minimum average brightness value for the lower cheek template 730 is lowered.

If the cheek brightness process 2020 passes, then the cheek-eye brightness ratio process 2025 determines if the ratio of the left and right cheek templates 728 and 730 to the respective eye templates 720 and 722 exceed a threshold value. This process accounts for the fact that the user's cheeks are brighter than his eyes. Process 2025 is described below with reference to Figure 18h. At step 2025a, the average gray level value for each of the eye templates 720 and 722 is calculated. At step 2025b, the average gray level value for each of the cheek templates 728 and 730 is calculated. Alternatively, step 2025b does not need to be performed if the average gray level calculated for the cheek templates at process 2020 is used. Next at step 2025c, it is determined if the average pixel value is greater than a threshold of, for example, 250. In very bright lighting, the cheek-eye ratio becomes smaller due to imager saturation. Thus, step 2025c determines if the lighting is very bright by determining if the average pixel value is above the threshold.

At steps 2025d and 2025e, if the lighting is bright, it is determined if the check-eye ratio value is greater than a CHEEKRATIO1 threshold of, for example, 1.25 for the right eye and cheek templates and greater than a CHEEKRATIO2 threshold of, for example, 1.1 for the left eye and cheek templates. At steps 2025f and 2025g, if the lighting is not bright, it is determined if the check-eye ratio value is greater than a CHEEKRATIO3 threshold of, for example, 1.4 for the right eye and cheek templates and greater than a CHEEKRATIO4 threshold of, for example, 1.2 for the left eye and cheek templates. If the check-eye ratio satisfies these criteria, process 2030 is performed. Otherwise, process 2010 is performed.

The bridge of the nose energy process 2030 determines if the average X-Laplacian energy (squared X-Laplacian value) of the mid-eye or bridge of the nose template 726 (shown in Figure 18d) is at most, for example, half of the X-Laplacian energy of the eye template 720 or 722 that has lower X-Laplacian energy (shown in Figure 18d). This process looks for a texture free region between the eyes. Hence, the bridge of the nose energy process 2030 is useful for eliminating false detection of eye pairs from the user's hair or other regions having high texture. As with the lower eye

template 722, the mid-eye template 726 may be varied. The bridge of the nose energy process 2030 is described below with reference to Figure 18i.

At step 2030a, the mid-eye template 726 is moved three pixels towards and away from the mouth template 732 and the average X-Laplacian energy is calculated for each position of the mid-eye template 726. At step 2030b, the mid-eye template that has the largest average X-Laplacian energy is selected. At step 2030c, the average X-Laplacian energy for each of the eye templates 720 and 722 are calculated. At step 2030d, if the largest average X-Laplacian energy of the selected mid-eye template is less than half the average X-Laplacian energy of the eye template 720 and 722 that has the lower average X-Laplacian energy of the two eye templates, then process 2035 is performed. Otherwise, process 2010 is implemented.

In Figure 18b, the bridge of the nose brightness process 2035 determines whether the brightness of each of the eye templates 720 and 722 (shown in Figure 18c) is no greater than, for example, .8 times the brightness of the bridge of the nose template 726 (shown in Figure 18c). The average gray level value is calculated for the bridge of the nose template 726 selected in step 2030b shown in Figure 18i. The average gray level value may also be calculated when the average X-Laplacian values are calculated at step 2030a shown in Figure 18i. The average gray level value is also calculated for each of the eye templates 720 and 722. Then, it is determined whether the average gray level value for each of the eye templates 720 and 722 is no greater than .8 times the average gray level value for the bridge of the nose template 726. If this criteria is satisfied, process 2040 is implemented. Otherwise, processing passes to process 2010.

The cheek energy process 2040 determines if the energy of level of the cheek templates is below a maximum. First, the Laplacian values for each of the cheek templates 728 and 720 (shown in Figure 18c) are computed and squared. Then the average of the squared values is computed and it is determined if the average value is below a threshold value of, for example, 45. In the exemplary embodiment, given the placement of the light sources, the threshold for the upper cheek template 728 is 40 and the threshold for the lower cheek template 730 is 80. The threshold for the lower cheek template 730 is higher than the threshold for the upper cheek template 728 because of the reduced illumination in the area of the lower cheek template 728. If this criteria is satisfied, process 2045 is implemented. Otherwise, processing passes to process 2010.

The mouth process 2045 is the same process 2085. If the criteria of the mouth process 2045 is satisfied, process 2050 is implemented. Otherwise, processing passes to process 2010. The nose process 2050 is the same process 2090. The eye position coordinates are stored at step 2055.

At step 2094, if any eye positions are stored at step 2055, steps 2060 through 2074 are implemented as described above. Otherwise, as shown in Figure 18v, the WFOV corneal specularities detector process 2096 is implemented. This is done because in certain situations, the template-based face finder may fail when, for example, the user is wearing a face mask covering the nose and the mouth. In such situations, the system falls back to an alternative scheme of WFOV eye-localization based on WFOV corneal specularities.

Process 2096 is the same as process 2071 described above. Next, at step 2098, the coordinates of the detected corneal specularities having the largest Y coordinate value is selected as the localization, the position of the user's eye. The corneal specularity having the largest Y coordinate value is most likely on the user's right eye.

Figure 28 shows another process for detecting eyes using reflection off the back of the eye and the occluding property of the iris/pupil boundary. The problem is to uniquely locate the position of an eye in both WFOV imagery and NFOV imagery. An example of WFOV imagery in this context is when the head of a person occupies approximately 1/3 of the width of the image. An example of NFOV imagery in this context is when the iris in the eye occupies approximately 1/3 of the width of the image.

The procedure combines two constraints. The first constraint is that the retina reflects light directed towards it. This is popularly known as the "red-eye" effect in visible light. The second constraint uses the geometry of the eye. Particularly, the occluding property of the iris-pupil boundary is used. The "red-eye" effect occurs when light is directed through the pupil and reflected off the retina and straight back in the same direction as the illumination into an observing imager. The three dimensional geometry of the eye is such that if the illumination is placed slightly off-axis compared to the imager, the light is reflected off the retina onto the back of the iris rather than back through the pupil. As a result, no or little light is returned. Small changes in illumination direction cause relatively large changes in the response from the eye. Further, the geometry of other features in the image is usually such that small changes in illumination direction has very little impact on the illumination returned to the imager. This property is used to differentiate bright points from the "red-eye" effect from bright points elsewhere.

At step 3000, shown in Figure 28, a first image F1 of a person is acquired using a imager with on-axis illumination 321a, shown in Figure 5c, turned on and the off-axis illumination 321b, shown in Figure 5c, turned off. At step 3010, a second image F2 of the person is recorded 1/30 second later with the on-axis illumination 321a turned off and the off-axis illumination 321b turned on. Next, at step 3015, the image F1 is multiplied by a constant and an offset is added to account for variations between the two images. At step 3020, the first image F1 is subtracted from the second Image

F2 to produce image F3. At step 3025, a circle detecting procedure is implemented to detect portions of the image which may be the eye. A spoke detecting algorithm may be used. The detected portions of the image are then processed in the remaining steps of the procedure. At step 3030, the absolute value of Image F3, the remaining portions, is generated. At step 3040, a Gaussian pyramid of the absolute difference image is generated to produce a Gaussian pyramid image. At step 3050, the values in the Gaussian pyramid image are compared to a threshold value. At step 3060, the regions of the Gaussian pyramid image above the threshold value are selected. At step 3070, the selected regions in the image above the threshold value are designated as regions corresponding to the user's eye.

When the images are subtracted at step 3020 the pupil effect is increased. The pupil effect is strong with on axis illumination. There may be clutter in the image after the subtraction. The clutter may be removed using a ring of LEDs in a circle. The ring is off axis and circular to correspond to the shape of the pupil. For example, off axis LEDs at 0 degrees and every 90 degrees can be used to turn off and on separately and to obtain a difference image for each LED. Because the pupil is circular a bright difference image at the pupil is obtained for each difference images. For objects such as glass frames, the difference images will only be bright at one particular angle because glass frames are usually linear at least at the scale of the pupil. Thus, an illuminator such as an IR illuminator is used at each different off axis point in combination with ambient illumination. Each image acquired with a respective IR illuminator is subtracted from the image with ambient illumination. This produces a set of difference images.

If the difference images are obtained using LEDs at 0 degrees, 90 degrees, 180 degrees, and 270 degrees, four difference images are produced and the pupil appears as a bright difference at every angle in the four difference images and clutter in the images varies between each image. At some images it might be visible and in others it would not be visible. The change in all of the difference images is identified as clutter and the pupil is the area that does not change.

In order to obtain a clear, sharp image of the user's iris for the iris classification and comparison process 326 (shown in Figure 3), it is desirable to have an accurate measurement of the distance between the NFOV imager and the user's eyes. If the Locate Head step 1512 (shown in Figure 15) generated a range map in order to locate the head, then this range information is already known. If this method was not used, then the eye location process, described above in Figures 15 - 18a can be used with a pair of stereoscopic images, provided by imagers 10 and 12, as shown in Figure 19 to generate an accurate distance measurement. This is an implementation of the Find Range step 1414 of Figure 14.

The first two steps in the process shown in Figure 19, steps 1910 and 1912, separately process the right and left WFOV images, as described above with respect to Figures 15, 16, 17 and 18a, to locate the user's eyes in each image. Next, at step 1914, the X, Y coordinate positions of each eye in the two images are determined by the stereo module 316, shown in Figure 3. Also at step 1914, the stereo module 316 calculates the angle from each imager to each of the eyes, based on the determined pairs of X, Y coordinates. Knowing these angles, the focal length of the lenses used in the WFOV imagers 10 and 12 and the distance between the imagers 10 and 12, the distance of each eye from each of the WFOV imagers can be calculated using simple trigonometry. Because the relative geometry of the WFOV imagers 10 and 12 and the NFOV imager 14 is known, this distance can easily be converted into a distance between each eye and the NFOV imager.

If the second imager 12 is not used but the user is alternately illuminated from the right and left by, for example, the light sources 126 and 128, a similar method may be used to determine the Z coordinate distance. This method differs from that outlined above in that the points being compared would not be the determined user eye positions but points corresponding to shadow boundaries (e.g. the shadows cast by the customer's nose) in the two images. Alternatively, the relative positions of the corresponding specular reflections in each of the eyes may be used to determine the Z coordinate distance between the NFOV imager and the user's eyes.

If none of the methods outlined above is used to determine the Z coordinate distance, the NFOV imager may still be able to obtain a sharp image of the eye using either conventional autofocus techniques or autofocus techniques that are tuned to characteristic features of the eye. In this instance, it may be desirable to provide the NFOV imager 14 with an approximate Z coordinate distance. This approximate distance may be provided by the sonic rangefinder 332, shown in Figure 3, as is well known.

Once the range or Z coordinate distance has been determined, the next step in the process shown in Figure 14, step 1416, is to locate the iris using the NFOV imager 14. A process for implementing this step is shown in Figure 20. The first step in this process, step 2010 is to adjust the mirror 16, shown in Figures 1, 2 and 3, to capture an image at the X, Y and at least approximate Z coordinates, determined by the preceding steps of Figure 14. When the mirror 16 has been positioned, the specular light sources 126 and 128 are turned on and step 2012 is executed to obtain a low-resolution NFOV image. This image has a resolution of 160 by 120 pixels. Next, at step 2014, this image is scanned for specularities to verify that the image contains an eye. If specularities are found, then the X, Y, Z coordinates of the center point of the specularities are marked as the new position of the eye. Next, at step 2016, the driver

320 changes the focus of the near field of view imager to determine if the image can be made sharper. Even for a imager 14 that has an autofocus feature, if the customer is wearing glasses, the imager 14 may focus on the glasses and may not be able to obtain a sharp image of the iris. Step 2016 compensates for corrective lenses and other
5 focusing problems by changing the focal length of the imager 14 and monitoring the image for well defined textures, characteristic of the iris.

One method for implementing autofocus is shown in Figure 27. As shown in Figure 27, at step 3605, the NFOV imager is adjusted using the adjustment values. At step 3610, an image is acquired and, at step 3615, the system tries to identify the
10 person's eye in the image. At step 3620, the process proceeds to step 3625 if it is determined that the person's eye was identified. If a person's eye is not identified in the image, the process proceeds to step 3625a.

At step 3625, additional images are acquired in front of and in back of the depth value. This is implemented by changing the focus of the NFOV imager 14 to a point
15 greater than the depth value. The NFOV imager 14 is then moved from this focus to a focus on a point which is closer to the NFOV imager 14 than the depth value. As the NFOV imager 14 is moved from the farthest point to the nearest point in the range, the system acquires images at periodic intervals. For example, five image images may be acquired over the entire range.

In an alternative embodiment, the focus of the NFOV imager 14 can be adjusted to a point in the 3-D space on either side of the depth value. An image can be acquired at this point. Then the focus of the image, as described below, can be obtained. If the focus of the newly acquired image is better than the previously acquired image, then the focus of the NFOV imager 14 is adjusted in the same direction to determine if a more
20 focused image can be acquired. If the newly acquired image is not in better focus than the previously acquired image, then the system proceeds in the opposite direction to determine if a better focused image can be obtained. The first embodiment is more advantageous than this embodiment because the delays required for readjusting the focus of the NFOV imager 14 can be avoided.

At step 3630, a first one of the acquired images are selected. A Laplacian pyramid of the image region containing the person's eye of the selected image is generated. At step 3640, the values in the Laplacian image L0 from the Laplacian pyramid are squared. Then, at step 3645 the squared values are summed. At step 3650 the summed values are divided by the values in Laplacian image L1. Prior to division,
35 the values in Laplacian image L1 are squared and summed. The summed values of Laplacian image L1 are used to divide the summed values of Laplacian image L0. The calculated values are stored in a memory (not shown). At step 3655, it is determined if each of the acquired images has been processed. If each image has not been processed,

at step 3660, the next one of the acquired images is selected and step 3635 is repeated. If all of the acquired images have been processed, the calculated values for each image are compared and, at step 3670, the image having the greatest calculated value is selected as the focused image.

5 The process performed by steps 3652A, 3630A, 3635A, 3640A, 3645A, 3650A, 3655A, and 3660A, are the same as the process described above with regard to steps 3625 to 3665 except that the image processing is performed over the entire obtained image because the person's eye was not identified. The processing performed in steps 3625 to 3665 is performed for the region which includes the user's eye.

10 Next, at step 2018, the iris preprocessor 324 determines whether specularities were found at step 2014. If specularities were not found, it may be because the image was not in focus. Accordingly, at step 2020, the process determines if step 2014 was executed before the image was focused. If so, then control is transferred to step 2014 to try once again to find specularities. If, after two tries, no specularities are found at
15 step 2014, step 2020 transfers control to step 2022 which attempts to find an eye in the image by locating eye-specific features. One possible implementation of step 2022 is to use the circle finder algorithm described below with reference to Figures 22a and 22b. Another implementation may be to search the low-resolution NFOV image for a dark central area surrounded by a brighter area in much the same way as described above with reference to Figure 17. Step 2022 selects the X, Y, Z coordinates for the eye from
20 the best candidate possible eye that it locates.

 After step 2022 or after step 2018 if specularities were found in the image, step 2024 is executed to correct the image for pan and tilt distortion. This step also centers the NFOV image at the determined X, Y coordinate position. The step 2024 warps the
25 near field of view image according to a rotational transformation in order to compensate for rotational distortion introduced by the mirror 16. Because the NFOV image is captured using a mirror which may be tilted relative to the imager, the image may exhibit rotational distortion. The iris preprocessor 324 may compensate for this rotational distortion by warping the image. Since the two tilt angles of the mirror 16 are
30 known, the warp needed to correct for this distortion is also known. While the warp transformations can be calculated mathematically, it may be simpler to empirically calibrate the warp transformations for the pan and tilt mirror prior to calibrating the X, Y, Z coordinate mapping between the WFOV image and the NFOV image.

 Figure 21 is a flow-chart of a process suitable for implementing the Obtain High
35 Quality Image step 1418 of Figure 14. As indicated by the dashed-line arrow 2115, a portion of this process is optional. The portion of the process bridged by the arrow 2115 performs much the same function as the focusing process described above with reference to Figure 20. This process, however, operates on the high-resolution image.

The focusing method illustrated by steps 2116, 2118 and 2120 may produce better results than a conventional autofocus algorithm for users who wear glasses. The focusing techniques described above may focus on the front of the glasses instead on eye features. The eye-specific autofocus algorithm, however, focuses on the eye even through many types of corrective lenses. This process may replace or augment the focusing process described above with reference to Figure 20.

The first step in this process, step 2110, adjusts the mirror 16 for the modified X, Y, Z coordinates, determined by the process shown in Figure 20. Next, at step 2112, the iris preprocessor 324 obtains a high-resolution image. In the exemplary embodiment, this is a 640 pixel by 480 pixel image. At step 2113, this image is scanned for specular reflections. The specular reflections are used both to confirm that the image contains an eye and, as described below, to determine gaze direction.

The specular reflections detected at step 2113 may be confirmed by using a number of checks. The first check is to calculate the difference in brightness or brightness ratio between the candidate point and a point at a distance D (e.g. 10 pixel positions) to the left, above and below the candidate point. If the candidate point is brighter by at least a threshold value than all of these points, then a search is performed for the second specularity. The search region can be determined by the physical separation of the lights, the distance between the customer and the lights and the nominal radius of curvature of the eye. If a second candidate specularity is found, then it is compared in brightness to its surrounding pixels. If this threshold test is passed then the specularity has been detected. The above description assumes that the specular reflections are generated by light sources which do not have any characteristic shape. If shaped light sources are used, such as the sources 126 and 128, shown in Figure 1c, an additional comparison may be made by correlating to the shapes of the reflections.

Figures 21a and 21b illustrate an alternative process for locating specularities in the NFOV imagery. Steps 8000 through 8045 are a process for locating a specularity in the NFOV imagery. At each step 8005, 8015, 8025, and 8035 a difference between a candidate pixel and a pixel ten pixels to the top, left, right or below of the candidate pixel is obtained. At steps 8010, 8020, 8030, and 8040, it is determined whether the candidate pixel is greater than 3 units of those pixels using the difference. If so, the candidate pixel is stored at step 8045 and a search is conducted to locate another specularity in a specified region to the right of the stored candidate pixel as determined at step 8060. A similar process is repeated for the pixels in the specified region in steps 8065 through 9093 as were performed in steps 8000 through 8045. Once candidate pixels have been obtained, the candidate pixels are verified in steps 8095 through 8099 to determine whether the stored candidate pixels are specularities. At step 8095, a difference is obtained between the candidate pixels and each pixel ten pixels away from

the candidate pixel as shown in Figure 21b. At step 8097, it is then determined whether each of the candidate pixel is greater than each of the pixels by three units using the difference. If so, at step 8099, it is determined whether the pixels which passed the above steps form at least a two by two pixel group of the candidate pixels that have passed the steps above to verify that the candidates pixels are specularities.

Next, at step 2114, the image is corrected for rotational distortion introduced by the mirror 16. At step 2116, the boundaries of the high-resolution image are stepped in the X and Y directions and at step 2118, the stepped image is scanned for sharp circular features. This step may be implemented using a circle finder algorithm, such as that described below with reference to Figures 22a and 22b and stepping the focus of the NFOV imager when circular features are found to determine if they can be made sharper. Alternatively, or in conjunction with the circle finder algorithm, step 2118 may change the focus of the near field of view imager 14 to achieve the best image textures. If the image was properly centered when the low-resolution near field of view image was processed, the main source of textures is the customer's iris. Until step 2120 indicates that a sharp pupil and/or textured iris have been found, step 2120 transfers control back to step 2116 to continue the search for the eye.

If the steps 2116, 2118 and 2120 are bypassed as indicated by the dashed-line arrow 2115, then the pupil and iris locations determined by the method described above with reference to Figure 20 are translated from the low-resolution NFOV image to the high-resolution NFOV image and used by the steps that follow. Otherwise the locations determined at steps 2118 and 2120 are used.

At step 2124, the position of the specular reflection found at step 2113 is compared to the position of the pupil, as determined as step 2118. If the X, Y, Z coordinate position of the eye is known, then the position of the specularity, relative to the pupil or iris boundary when the customer is looking directly at the NFOV imager 14 is also known. Any displacement from this position indicates that the customer is not looking at the NFOV imager and, thus, that the image of the iris may be rotationally warped compared to the image that would be most desirable to pass to the classification and comparison process 326 (shown in Figure 3). The warping needed to correct the image of the iris is uniquely determined by a vector between the desired position of the specular reflection and its actual position in the image. At step 2124, this vector is determined and the image is warped to correct for any rotational distortion caused by the customer not looking directly at the NFOV imager. The type and magnitude of this warping can be determined in the same way as the warping used to correct for rotational distortion caused by the mirror 16.

Alternatively, any rotational distortion of the image caused by gaze direction can be determined directly from the shape of the iris. If the iris is determined to be oval the

correction needed to make it circular can be determined and applied using well-known techniques.

The next step in the process, step 2126, determines if the specularity in the NFOV image obscures too great a portion of the iris for accurate recognition. If so, it may be desirable to obtain a new high resolution image with the light sources 126 and/or 128 turned off. To ensure proper alignment, this image may be captured as a matter of course immediately after the image which includes the specular reflection but not analyzed until it is determined that is desirable to do so.

If a new image is obtained at step 2128, step 2130 is executed to correct for the pan and tilt and gaze direction rotational distortions. Steps 2126, 2128, and 2130 are performed iteratively. Once these steps have been executed, the normalized image may be processed once again using the circle finder algorithm to locate the iris boundaries. These include the limbic boundary between the iris and the cornea, the pupil boundary between the iris and the pupil and the eyelid boundaries which may obscure portions of the top and bottom of the iris. details of the process used to implement this step are described below with reference to Figure 23. At step 2133, once the iris boundaries in the image have been determined, that portion of the image corresponding only to the customer's iris are extracted and passed to the iris classification and comparison process 326.

Figure 22a is a flow-chart illustrating an implementation of the circle finder process used in the processes shown in Figures 20 and 21. Briefly, this process sequentially selects points in the image as target center points of the eye. It then determines a cost function for edge information of image pixels lying along spokes emanating from the target center point. The center point having the lowest cost function is designated as the center point of the iris.

The first step in the process, step 2210, is to locate edges in the image of the eye. This step may use any of a number of edge finder algorithms. In the exemplary embodiment, the image of the eye is first low-pass filtered to reduce noise using, for example, the Gaussian image produced by a pyramid processor. Next, a three-tap FIR filter, having weighting coefficients (-1, 0 and 1) is used to locate vertical edges by scanning the horizontal lines through the filter. As a last edge detection step, the same filter is applied to the vertical columns of the image to locate horizontal edges.

After step 2110, step 2212 and the steps which follow it select a pixel that is not an edge pixel and calculate a cost function for edges located along 12 spokes emanating from the selected point. This operation is illustrated in Figure 22b. In Figure 22b, the selected center point is the point 2240. Twelve spokes 2250 are defined as emanating from that point. The spokes are biased in the horizontal direction because the circular boundaries of the iris in the vertical direction may be occluded by the eyelid. The

spokes are separated by a ten degree angle. Due to the difference in luminance level, it is expected that the iris-sclera boundary is characterized by relatively well-defined edges. The starting and ending points of the spokes are selected using prior knowledge of the expected radius of the iris.

5 At step 2214, along each spoke, the edge having a magnitude that is greater than a threshold value and an orientation that matches most closely the predicted edge orientation (given by the spoke angle) is selected as the candidate iris/sclera boundary. This step is performed for each spoke at step 2216. Then, at step 2218, for all of the edges, the median distance from the candidate center point is calculated. Also at step
10 2218, selected edges that have a distance much greater than this calculated distance are discarded as being outliers caused, for example, by specularities in the eye. At step 2220, a cost is calculated for the remaining edges. This cost is the sum of the absolute difference between the predicted edge orientation and the measured orientation, multiplied by a normalization factor and added to the sum of the absolute difference
15 between the median radius and the measured radius, multiplied by a normalization factor. At step 2222 this process is repeated for all candidate center pixels. For a perfect circle, this cost is zero.

 At step 2224, the process determines whether any of the possible center points has a cost value that is less than a threshold value. This threshold value may be
20 determined objectively as a low expected edge strength value for the iris-sclera boundary or it may be determined subjectively from data that was captured for the customer when her iris was first registered with the system. If no cost value is found to be below the threshold value, the process is unsuccessful in finding the iris. If the circle finder is invoked at step 2118 of Figure 21, this is an expected possible outcome.
25 If it is invoked at step 2132, however, the "not found" response indicates that the iris location process has been unsuccessful. In this instance, the control process 310 may attempt to retry the process or may use other means of verifying the identity of the customer.

 At step 2226, the minimum cost of any of the candidate center pixels is selected
30 if more than one cost value is below a predetermined threshold. Once a candidate center point is determined, then at step 2228, the iris-sclera boundary is defined for the image, at step 2230, the pixels lying along the spokes are analyzed once again to detect the iris-pupil boundary and that boundary is defined at step 2232. After step 2232, step 2234 is executed to notify the process which invoked the circle finder algorithm that an
35 eye had been successfully located.

 The last step in extracting the iris is to eliminate pixels corresponding to the eyelids. A process for implementing this step is shown in Figure 23. The first step in this process, step 2310, identifies the image area between the iris-sclera boundary and

the iris-pupil boundary. Next, at step 2312, the high-resolution image outside of the iris-sclera boundary and inside the iris-pupil boundary are blanked. This image is then analyzed at step 2314 for horizontal and near horizontal edges. An exemplary algorithm for finding edges with a particular orientation is Fisher's Linear Discriminate which is described in section 4.10 of a textbook by R.O. Duda et al. entitled *Pattern Classification and Scene Analysis*, Wiley Interscience, 1974, pp. 114-18. As described above, this step is optional, depending on the iris recognition algorithm that is used.

If, at step 2316, horizontal or near-horizontal edges are found in the partially blanked image then, at step 2318, the image is blanked above and below the edges. This blanked image includes only components of the customer's iris. It is this image which is passed to the classification and comparison process 326.

In the exemplary embodiment, the classification and comparison process 326 is the process described in U.S. Patents Nos. 4,641,349 and 5,291,560, which are hereby incorporated by reference for their teachings on iris recognition systems. It is contemplated, however that other iris recognition systems could be used. One exemplary system may be implemented using spatial subband filters. These filters would be applied to a central band of the extracted iris image to generate multiple spatial frequency spectra, each corresponding to a separate subband. A comparison with the image would include generating a similar set of subband frequency spectra for the customer's eye from the high-resolution image and comparing the various frequency spectra to determine the likelihood that the iris being imaged has the same characteristics as are stored in the customer database. Alternatively or in addition, respective subband images of the scanned iris and the stored may be correlated to determine the likelihood of a match.

The WFOV/NFOV recognition system described above is not limited to iris recognition but may be applied more generally to any system in which an object having both large-scale and small-scale features is to be identified. One such application is illustrated in Figure 33. In this application, three recognition systems, 9710, 9720 and 9730, such as the one described above are used in a warehouse for inventory control. In the system shown in Figure 33, a worker is removing a box 9716 from the warehouse using a forklift 9718. The three recognition systems continually scan the scene for large-scale features indicative of barcodes. When such a large-scale feature is identified by the WFOV processing, the system switches to NFOV processing to capture and analyze an image of the barcode.

While the barcodes may be applied to materials stored in the warehouse, they also may be applied to equipment, such as the forklift 9718 and to the hardhat 9719 worn by the worker. Thus, the recognition system can be the sensing system of an

inventory control and management system for the warehouse. In addition, matching barcodes may be applied to all sides of the materials in the warehouse and, as each barcode is read by the recognition system, it is checked against barcodes read by the other systems for consistency.

5 As shown in Figure 34, as the forklift approaches the loading dock, the recognition system 9720 scans the scene for an identifying bar code, such as the barcode 9810. An exemplary barcode that may be used on the box is shown in Figure 35B while a perspective view of the box 9716 is shown in Figure 35A. As shown, the barcode is composed of concentric circles and is printed on all sides of the box 9716. It
10 may also be printed with a reflective or shiny surface which would provide specular reflections. Accordingly, the features of the system described above which locate a target object in the WFOV image using specular reflections and then use image processing techniques based on circular features to extract a region of interest from the NFOV image may be directly applied to obtaining an image of the barcode 9910. In
15 addition, the circle finding algorithm described above with reference to Figures 20 and 21 may be used both to locate the barcode in the NFOV image and to "read" the barcode by analyzing the pattern of light and dark regions along the most likely diameter.

The system may use a more conventional linear barcode 9916 as shown in Figure 35C. This barcode may also be printed with a reflective surface or it may
20 include markers such as 9914 which emit light having a predetermined frequency when excited, for example, by ultraviolet light. These markers may be used in much the same way as the "red-eye" effect described above, either to locate the barcode in the WFOV image or to determine the proper orientation of the image in the NFOV image. Instead of placing the markers at the bottom of the label it is contemplated that the
25 markers may be placed on either side of the linear barcode, indicating a path to be followed to read the barcode.

It is also contemplated that the infrared imaging techniques described above may be used to compensate for low-light conditions in the warehouse or to provide added security by having a barcode that is visible only in infrared light printed at some known
30 offset from the visible light barcode. The near field of view image may be aligned to capture and process both the visible and infrared barcode images.

Another recognition task to which the exemplary system may be well suited is capturing images of license plates on vehicles. The location of a license plate may vary greatly with the type of vehicle. License plates, however, generally conform to a
35 known size and shape. Used for this purpose, the WFOV processing would search images for features that are likely to be classified as license plates and then the NFOV processing may focus on the area of the scene identified by the WFOV imager and isolate a region of interest that may contain a license plate. In this application, the

image recognition system may use such features as the reflectivity of the license plate or the infrared signature of tailpipe emissions to isolate an area of interest in the WFOV image.

5 A system of this type may be useful for monitoring vehicles entering and leaving a secure installation. Furthermore, by adding image stabilizers (not shown) to both the WFOV and NFOV processors, the system may be used in police vehicles to obtain an image of the license plate on a vehicle.

10 In both of the systems described above, it may be desirable to use a multi-level rangefinding technique which may, for example, determine a rough distance to the object to be imaged from an acoustic sensor, use this rough distance to focus the WFOV imager and to determine a zoom distance for the NFOV imager and then use other techniques, such as the stereoscopic range finding technique described above to provide coarse region of interest information and focus information to the NFOV processor. With this information, the NFOV processing can capture a focused image
15 and refine the region of interest to obtain a detailed focused image of the target, either the barcode or the license plate.

20 The invention is also a method for obtaining and analyzing images of at least one object in a scene comprising the steps of capturing a wide field of view image of the object to locate the object in the scene and using a narrow field of view imager responsive to the location information provided in the capturing step to obtain higher resolution image of the object.

25 The invention is also a method for obtaining and analyzing images of an object in a scene comprising the steps of capturing an image of the scene, processing the image at a coarse resolution to locate the object in a region of interest in the image, and processing the region of interest at a second resolution greater than the first resolution to capture a high resolution image of the object.

While the invention has been described in terms of multiple exemplary embodiments, it is contemplated that it may be practiced as outlined above within the spirit and scope of the following claims.

The Invention Claimed is:

1. An iris recognition system which obtains and analyzes images of an iris of at least one eye in a scene comprising:

a wide field of view imager which is used to capture an image of the scene and to locate the eye; and

a narrow field of view imager which is responsive to the location information provided by the wide field of view imager and which is used to capture an image of the iris of the eye, the image of the iris having a higher resolution than the image captured by the wide field of view imager.

2. The system of claim 1 in which the eye is located on a head and the system includes apparatus, coupled to the wide field of view imager and to a further wide field of view imager, which generates a depth map of an image from stereoscopic information provided by the wide field of view imager and the further wide field of view imager;

means, responsive to the depth map, for locating a head in the image; and
means, responsive to the determined location of the head in the image for locating the eyes.

3. The system of claim 1 in which the eye is located on a head and the system includes apparatus, coupled to the wide field of view imager which is responsive to motion in the image to locate the head in the image; and

means, responsive to the determined location of the head in the image for locating the eyes.

4. The system of claim 1 in which the eye is located on a head and the system includes apparatus, coupled to the wide field of view imager which is responsive to flesh tones in the image to locate the head in the image; and

means, responsive to the determined location of the head in the image for locating the eyes.

5. The system of claim 1 in which the eye is located on a head and the system includes:

apparatus, coupled to the wide field of view imager which is responsive to patterns in the image corresponding to a face to locate the head in the image; and

means, responsive to the determined location of the head in the image for locating the eyes.

6. The system of claim 1 in which the eye is located on a head and the system includes apparatus, coupled to the wide field of view imager which matches sections of the image corresponding to a face to a template to locate the eye in the image.

7. A system which obtains and analyzes images of at least one object in a scene comprising:

a wide field of view imager which is used to capture an image of the scene and to locate the object; and

a narrow field of view imager which is responsive to the location information provided by the wide field of view imager and which is used to capture an image of the object, the image of the object having a higher resolution than the image captured by the wide field of view imager.

8. The system of claim 7 wherein the object has a feature and the narrow field of view imager captures an image of the feature.

9. The system of claim 8 wherein the image of the feature excludes other parts of the object.

10. The system of claim 7 further comprising

a further wide field of view imager which is used to capture a further image of the scene;

means for producing a depth map from stereoscopic information provided by the wide field of view imager and the further wide field of view imager;

means, responsive to the depth map, for locating the object in the image; and

means, responsive to the determined location of the object in the image for locating a feature on the object.

11. A recognition system which obtains and analyzes images of at least one object in a scene comprising:

at least one wide field of view imager;

wide field of view circuitry, coupled to the wide field of view imager, which processes images provided from the wide field of view imager;

a narrow field of view imager coupled to the wide field of view circuitry; and

narrow field of view circuitry, coupled to the narrow field of view imager and to the wide field of view circuitry, which processes images provided from the narrow field of view imager.

12. The system of claim 11 wherein the wide field of view circuitry comprises:

control circuitry which controls operations of the narrow field of view imager and the wide field of view imager;

a crosspoint switch; and

at least one pyramid processor coupled to the crosspoint switch, the pyramid processor processes images provided from the wide field of view imager via the crosspoint switch.

13. The system of claim 12 wherein the narrow field of view circuitry comprises a pyramid processor, coupled to the crosspoint switch, that processes the images provided from the narrow field of view imager via the crosspoint switch.

14. A system which obtains and analyzes images of an object in a scene comprising:

means for capturing an image of the scene and processing the image at a coarse resolution to locate the object in a region of interest in the image; and

5 means for processing the region of interest at a second resolution greater than the first resolution to capture a high resolution image of the object.

15. A method for obtaining and analyzing images of at least one object in a scene comprising the steps of

10 capturing a wide field of view image of the object to locate the object in the scene; and

using a narrow field of view imager responsive to the location information provided in the capturing step to obtain higher resolution image of the object.

16. A method for obtaining and analyzing images of an object in a scene comprising:

15 capturing an image of the scene;

processing the image at a coarse resolution to locate the object in a region of interest in the image; and

processing the region of interest at a second resolution greater than the first resolution to capture a high resolution image of the object.

1/41




FIG. 1a

SUBSTITUTE SHEET (RULE 26)

2/41






FIG. 1b

SUBSTITUTE SHEET (RULE 26)

3/41

**FIG. 1c****FIG. 2**
SUBSTITUTE SHEET (RULE 26)

4/41




FIG. 3
SUBSTITUTE SHEET (RULE 26)

5/41




FIG. 4a




FIG. 4b




FIG. 5a




FIG. 5b

SUBSTITUTE SHEET (RULE 26)

6/41




FIG. 5c




FIG. 5d




FIG. 5e

SUBSTITUTE SHEET (RULE 26)

7/41




FIG. 6




FIG. 7




FIG. 8




FIG. 9

SUBSTITUTE SHEET (RULE 26)

8/41




FIG. 10




FIG. 11




FIG. 12




FIG. 13

9/41




FIG. 14

SUBSTITUTE SHEET (RULE 26)

10/41






FIG. 15

**FIG. 16**


SUBSTITUTE SHEET (RULE 26)

12/41

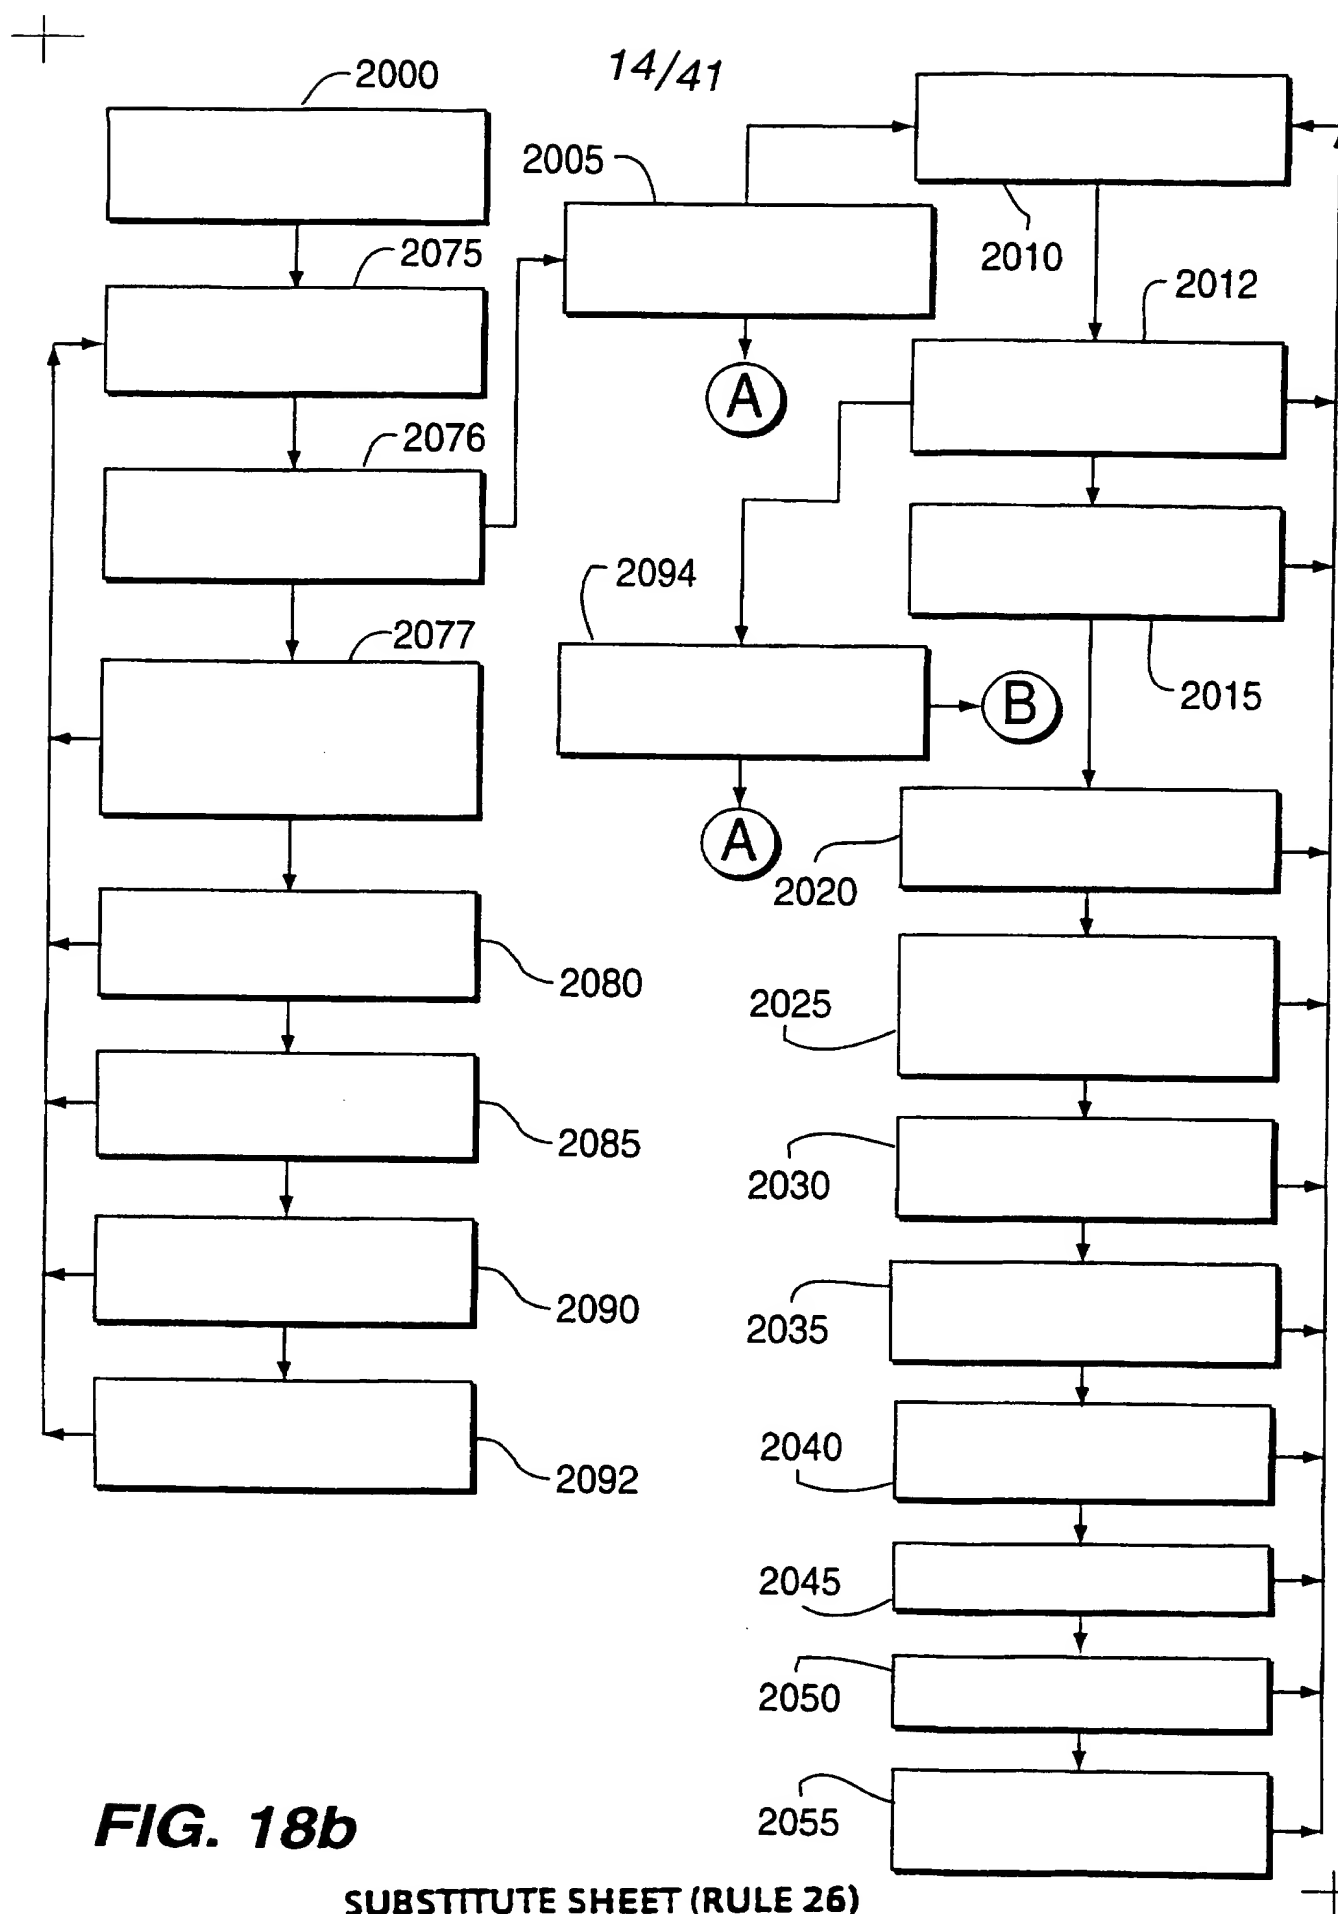
**FIG. 17**

SUBSTITUTE SHEET (RULE 26)

13/41

**FIG. 18a**

SUBSTITUTE SHEET (RULE 26)

**FIG. 18b**

SUBSTITUTE SHEET (RULE 26)

15/41




FIG. 18c

SUBSTITUTE SHEET (RULE 26)

16/41




FIG. 18d





FIG. 18e

SUBSTITUTE SHEET (RULE 26)

17/41

**FIG. 18f**

18/41

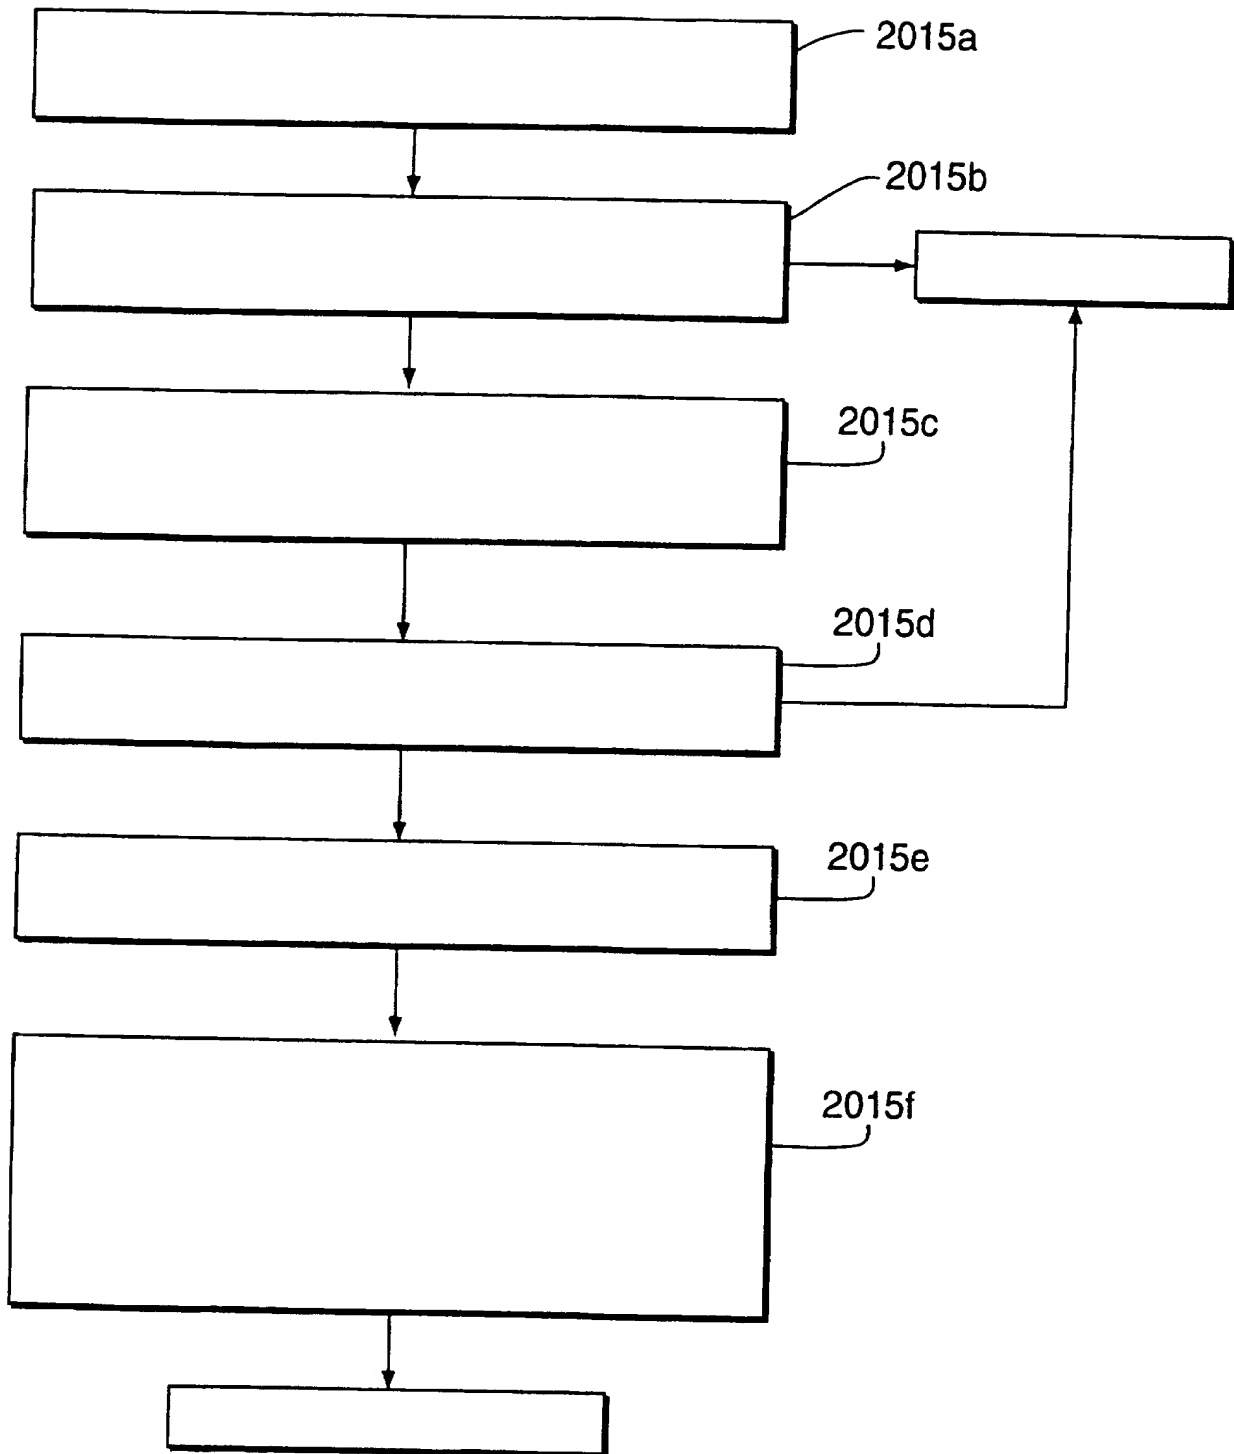


FIG. 18g
SUBSTITUTE SHEET (RULE 26)

19/41




FIG. 18h
SUBSTITUTE SHEET (RULE 26)

20/41




FIG. 18i
SUBSTITUTE SHEET (RULE 26)

21/41




FIG. 18j

SUBSTITUTE SHEET (RULE 26)

22/41




FIG. 18k






FIG. 18l

23/41



24/41




FIG. 18o




FIG. 18p




FIG. 18q




FIG. 18s




FIG. 18t




FIG. 18u

SUBSTITUTE SHEET (RULE 26)

25/41




FIG. 18r




FIG. 18v

26/41




FIG. 19





FIG. 20

27/41

**FIG. 21**

28/41





FIG. 21a

FIG. 21b

29/41

**FIG. 22a**

SUBSTITUTE SHEET (RULE 26)




FIG. 23





FIG. 22b

FIG. 24

32/41




FIG. 24a
SUBSTITUTE SHEET (RULE 26)

33/41







FIG. 24b

34/41

**FIG. 25****FIG. 26**

35/41

**FIG. 27**

36/41

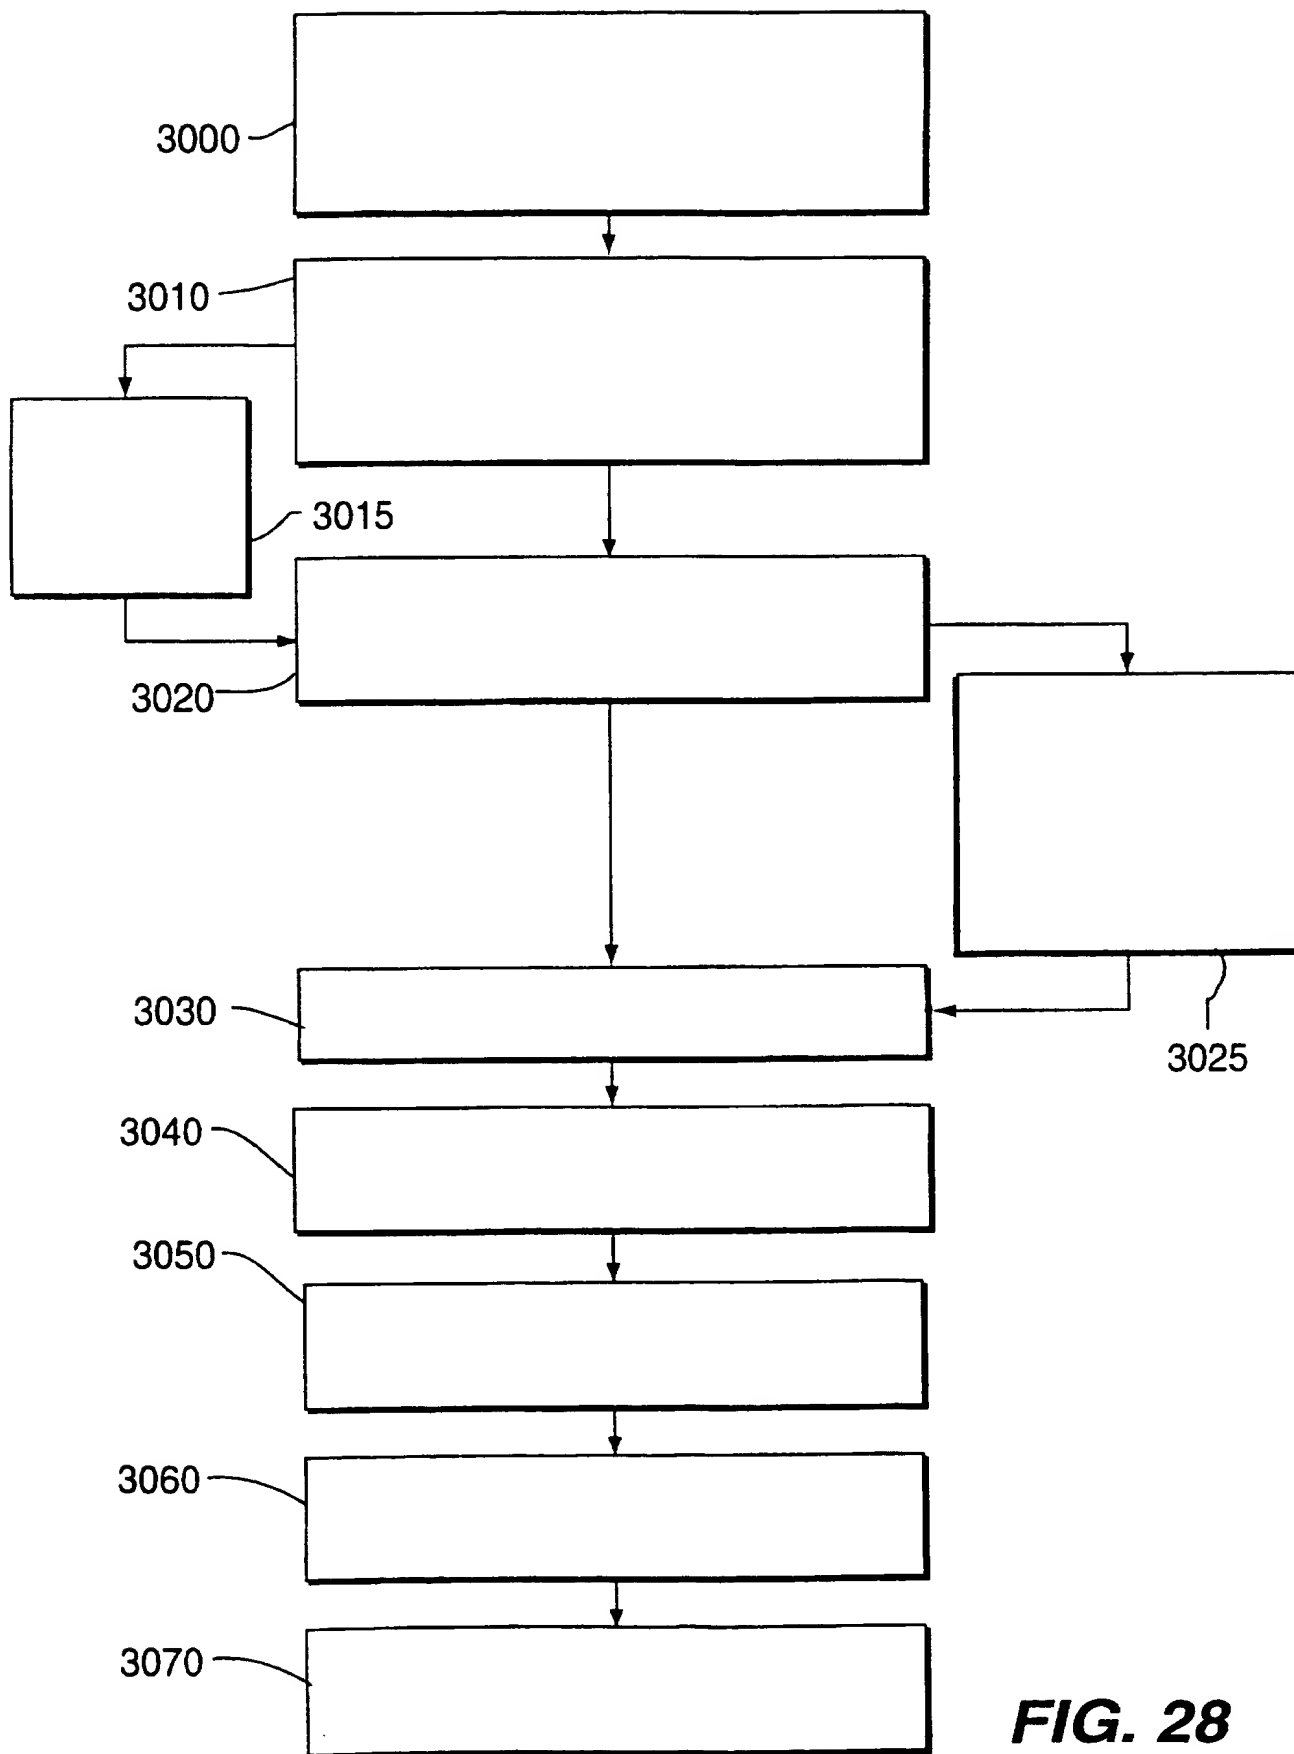




FIG. 28

37/41

**FIG. 29**

SUBSTITUTE SHEET (RULE 26)

38/41



SUBSTITUTE SHEET (RULE 26)

39/41




FIG. 31
SUBSTITUTE SHEET (RULE 26)

40/41




FIG. 32




FIG. 33

SUBSTITUTE SHEET (RULE 26)

41/41




FIG. 34

FIG. 35a





FIG. 35c

SUBSTITUTE SHEET (RULE 26)




FIG. 35b

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/19132**A. CLASSIFICATION OF SUBJECT MATTER**

IPC(6) : G06K 9/00

US CL : 382/117

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 382/115, 117, 128, 154, 240, 318, 319

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X,P ---- Y,P	US, A, 5,572,596 (WILDES ET AL.) 05 November 1996, Abstract; Figs. 1-3; column 2, lines 15-62; and column 3, line 12 to column 9, line 50.	1, 3-11, 14-16 ---- 2, 12, 13
Y	US, A, 5,063,603 (BURT) 05 November 1991, Abstract, Figs. 1-4, and column 2, line 23 to column 3, line 19.	2
Y	US, A, 5,452,376 (ESCALANTE RAMIREZ ET AL.) 18 September 1995, Abstract; Figs. 1-4, column 1, line 20 to column 2, line 59.	12, 13
Y	US, A, 5,329,599 (CURRY ET AL) 12 July 1994, Abstract, column 2, line 66 to column 3, line 50.	12, 13

☒ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	*T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
*A document defining the general state of the art which is not considered to be of particular relevance	*X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
*E earlier document published on or after the international filing date	*Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
*L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*G document member of the same patent family
*O document referring to an oral disclosure, use, exhibition or other means	
*P document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search

03 FEBRUARY 1997

Date of mailing of the international search report

07 MAR 1997

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

JOSEPH MANCUSO

Telephone No. (703) 305-8576

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/19132

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US, A, 5,430,809 (TOMITAKA) 04 July 1995, Abstract, Figs. 1-8, column 1, line 46 to column 2, line 37.	1-16
X	US, A, 4,589,140 (BISHOP ET AL) 13 May 1986, Abstract, Figs. 2, 4-12, column 1, line 8 to column 3, line 30.	1-16

Form PCT/ISA/210 (continuation of second sheet)(July 1992)*

THIS PAGE BLANK (USPTO)